HP Virtual Connect for the Cisco Network Administrator

"A technical discussion of the HP Virtual Connect 4.X features and their interoperation with a Cisco network infrastructure. Written from the perspective of a Networking person.

Abstract

This whitepaper discusses the networking technology behind Virtual Connect Ethernet as it relates to interoperability with a Cisco and/or other network infrastructures. This whitepaper specifically addresses the core Ethernet networking technologies used by Virtual Connect Ethernet and how they interoperate with a Cisco network infrastructure.

Disclaimer, Prerequisites, and Versioning

The purpose of this whitepaper is to provide a technical discussion of the Virtual Connect components, and their operation, with specific attention given to interoperability with Cisco switches and routers. This paper does not provide an exhaustive discussion of all aspects of Virtual Connect operation and configuration and this paper is not a competitive analysis between Virtual Connect and other vendor's Ethernet solutions.

It is assumed that the reader is already familiar with Cisco terminology and device operation and that the reader is familiar with the basics of HP BladeSystem c-Class enclosures, HP BladeSystem c- Class blade servers, and HP BladeSystem Virtual Connect. For additional information on these HP BladeSystem c-Class components, please visit:

- **HP Bladesystem**
- **HP Bladesystem Interconnects**
- **Beginnners Guide to HP Bladesystem**

Recommended Prerequisite Reading:

- **Overview to Virtual Connect Technologies.**
- **Virtual Connect User Guide (version 4.x)**

This whitepaper was written based on the features provided in Virtual Connect firmware version 4.x and earlier. Newer releases of firmware may introduce new features or may introduce changes to the way existing features work. For any discrepancies between the information in this paper and actual operation, it is recommended that the Administrator refer to the Virtual Connect manuals and release notes matching the firmware version being used. Both can be found online at www.hp.com/go/bladesystem/documentation under the "Install Your Solution tab".

Contents

Abstract 1
Disclaimer, Prerequisites, and Versioning 1
Introduction 5
  LAN-safe 5
  SAN-safe 5
Ethernet components and technology. 6
  Definitions 6
HP BladeSystem Architecture Overview 8
  c-Class Cabling Layout 11
  FlexNIC capabilities 12
HP Virtual Connect Technology 15
  Overview 15
Virtual Connect Components 16
  Virtual Connect Ethernet Modules: 16
    Virtual Connect Flex-10/10D Module 16
    Virtual Connect FlexFabric Module 17
  Virtual Connect Fibre Channel Modules 18
    Virtual Connect 24 Port Modules 18
    Virtual Connect 20 Port Module 18
    Virtual Connect Manager 19
Virtual Connect Manager User Interfaces 19
  VC MIB Support 19
  Virtual Connect Port Descriptions 20
  Overview of Virtual Connect Ethernet Connectivity 20
    VC Uplink Fault Tolerance 22
    VC Uplink Load Balancing 23
    Port Channeling (802.3ad) Load Balancing Algorithm 25
    VC Uplinks and VLAN Trunking 27
    Unsupported Cisco VLAN Protocols 29
    Virtual Connect and NIC Teaming 29
  Virtual Connect and Cisco Terminology Comparison 30
Sample Virtual Connect Ethernet and Cisco Nexus Configurations 30
Comparing VC and VMware Networking Technology 35
  Virtual Connect's Loop Prevention Technology 38
  Stacking Virtual Connect Ethernet Modules 39
Optimizing Virtual Connect Stacking Links 43
VC Managed MAC Addresses 43
Spanning Tree 45
Unidirectional Link Detection (UDLD) 45
Quality of Service 45
Security 46
  ACLs & VLAN ACLs 46
  Port Security 41
Private VLANs 42
Multicast & IGMP Snooping 43
SPAN (Port Monitoring) 43

Logging **44**
Centralized User Access Control (LDAP) **44**
Cisco Discovery Protocol (CDP) **44**
Additional Resources **45**
Appendixes **46**
Appendix A: Description of VC Network Diagram **46**
Appendix B: c3000 Port Mapping Diagram **47**
Appendix C: Frequently Asked Questions **48**

Introduction

Virtual Connect is an innovative networking product for HP bladesystem c-Class customers that has been shipping since February 2007. Virtual Connect was designed and engineered as a direct result of customer requests for a better way to manage blade server network connections. It is very important for the Administrator to fully understand the impact of the product and for them to feel comfortable with its introduction into their environment. Virtual Connects implementation into a customer environment will not only positively affect the Server Administrators, but also the LAN and SAN Administrators. While some will call it a switch, but it truly is a layer 1.5 device. It uses a mapping technology to create the data paths. Without the data paths being mapped data cannot flow. More specifically, because of this impact on LAN Administrators, a thorough discussion of Virtual Connect and its operation and interoperability with the existing Cisco network infrastructure is very important.

LAN-safe

Virtual Connect works seamlessly with your external network:

- Does **not** participate in Spanning Tree Protocol (STP) on the network uplinks to the data center. This avoids potential STP configuration errors that can negatively affect switches in the network and the servers connected to those switches.
- Uses an internal loop prevention algorithm to automatically detect and prevent loops inside the Virtual Connect domain. For example, if a Server Admin was to accidentally turn on Bridging at the OS layer, Virtual Connect will detect that and block one side of the loop and send a trap notifying the admins of this condition.
- Allows aggregation of uplinks to data center networks (using LACP and fail-over).
- Supports VLAN tagging on egress or pass-thru of VLAN tags in tunneled mode.
- Supports Link Layer Discovery Protocol (LLDP) and Jumbo Frames.
- Virtual Connect is not a Black Box. I some cases it provides more industry standard visibility than our competitors.
- Separate Data Plane and Control / Management Planes – True out-of-band management
- Authentication methods that include local authentication, Lightweight Directory Access Protocol (LDAP), Terminal Access Controller Access-Control System Plus (TACACS+), and Remote Authentication Dial-In User Service (RADIUS).
- Network Access Groups that let you control which networks to allow into the same server profile. You can assign a VLAN to one or more groups. This prevents administrators from using networks from different security domains in the same server profile.
- Virtual Connect supports multiple QOS methodologies.

SAN-safe

Virtual Connect Fibre Channel uplinks appear to be multiple HBAs connecting to the SAN by using N_Port ID Virtualization (NPIV) technology. NPIV is an industry-standard Fibre Channel protocol that provides a method to assign multiple Fibre Channel addresses on a single physical link. Each Fibre Channel connection has its own N_Port ID and WWN.

Virtual Connect works seamlessly with the external storage fabrics:

- Supports industry standard NPIV on both uplinks and downlinks.
- Doesn't consume Fibre Channel Domain IDs; therefore Virtual Connect doesn't affect the total number of devices that you can connect to an individual SAN Fabric.
- Compliant and compatible with SAN switches from any standards-based vendor.
- Transparently supports Cisco virtual storage area network (VSAN), Cisco inter VSAN routing (IVR), and Brocade Virtual Fabric features.
- Separate Data Plane and Control / Management Planes – True out-of-band management
- Authentication methods that include local authentication, Lightweight Directory Access Protocol (LDAP), Terminal Access Controller Access-Control System Plus (TACACS+), and Remote Authentication Dial-In User Service (RADIUS).

Virtual Connect modules supporting Fibre Channel must attach to NPIV-capable SAN switches. Most enterprise class SAN switches today support NPIV. It is possible to use NPV switches as an intermediary devices as long as it attaches to a NPIV device directly up stream. You can also connect VC FlexFabric modules directly to HP 3PAR StoreServ System arrays using the Virtual Connect 3.70 firmware (see the Direct-Attach Fibre Channel for 3PAR Storage Systems section).

Depending on the module, Virtual Connect-Fibre Channel modules can aggregate up to 255 physical or virtual server HBA ports through each of the module's uplink ports. This aggregation method is especially important to SAN administrators who struggle with SAN fabric segmentation and Fibre Channel Domain ID consumption.

Virtual Connect Fibre Channel modules make it easier to provision virtual machines by facilitating multiple HBA WWNs on the physical server. Each virtual machine can have its own unique WWN that remains associated with that virtual machine even when you move the virtual machine. Now SAN administrators can manage and provision storage to virtual HBAs, up to 128 per server blade, with the same methods and quality of service as physical HBAs.

Ethernet components and technology.

While this whitepaper assumes the reader is familiar with the basics of the HP BladeSystem c-Class architecture, it will review some of the important aspects that are key to understanding Virtual Connect technology/terminology. It will begin with a list of definitions for terms that will be used extensively within the paper. Next, a brief overview of the core networking components of the HP BladeSystem c- Class architecture will be covered, followed by an in-depth discussion of the Virtual Connect.

The subsequent sections will be devoted to covering the details of many networking features where Virtual Connect and the external Cisco infrastructure intersect.

Definitions

ACL (VACL)	Access Control List or VLAN Access Control List: A set of rules that allows or disallows network traffic to flow between network devices.

BPDU	Bridge Protocol Data Unit: A spanning tree configuration frame exchanged between switches in the same spanning tree domain
CDP	Cisco Discovery Protocol: A proprietary Cisco protocol used to exchange neighbor information between two directly connected Cisco devices
CX-4	An industry standard cabling specification used by VC for network connectivity using 10 Gbit Ethernet over copper.
DCBX	DCBX is a discovery and capability exchange protocol to discover peers and exchange configuration information between DCB–compliant bridges.
ETS	Enhanced transmission selection (ETS; IEEE 802.1 Qaz) Enables bandwidth management between traffic types for multiprotocol links.
External Network	The network and associated network devices external to the VC domain
FCOE	Fibre Channel Over Ethernet – Fibre Channel frames are encapsulated in a Ethernet Frame and transported over the Ethernet network.
Internal cross-connect	A non-visible port that interconnects two horizontally adjacent VC-Enet Modules – used for updating the Virtual Connect Managers on each of the Virtual Connect Modules.
LACP	Link Aggregation Control Protocol: An 802.1ax Link Aggregation configuration frame exchanged between two devices that form a port trunk\channel between them. This was originally part of the 802.3AD specification but was moved to the 802.ax – 2008 specification.
LAG	Link Aggregation Group. 802.1ax terminology for a port trunk\channel group
LLDP	Link Layer Discovery Protocol. An IEEE protocol that provides CDP-like functionality
Logical Path	A single physical port or a single port channel. Both represent a single communication path.
LOM	LAN on Motherboard. A NIC embedded on the system board of a server.
Oneview	HP's new Server Management platform.
PFC	Priority-based Flow control (PFC; IEEE 802.1 Qbb) Provides capability to manage bursty, single traffic source on a multiprotocol link.
Port Trunk (channel group)	A group of two or more ports that operate as a single logical port and single logical path for the purposes of load balancing. 802.3ad and EtherChannel are both port trunking technologies
Quality of Service (QoS)	A very broad term associated with network traffic classification, prioritization, queuing, marking, etc. Quality of service is particularly important for the transport of traffic with special requirements. In particular, much technology has been developed to allow computer networks to become as useful as telephone networks for audio conversations, as well as supporting new applications with even stricter service demands.
Server Profile	An object within the Virtual Connect domain that is assigned to a server bay

	and contains the server's LAN and SAN connectivity settings (vNet assignments, managed MAC addresses & WWNs, server boot parameters, PXE configuration, and fiber channel boot parameters).
SFP	A hot-pluggable modular 1 Gbit port. Pluggable modules allow for electrical or optical connectivity at 1 Gbit speeds
Shared Uplink Set (SUS)	The term used by Virtual Connect to configure one or more VC uplinks as a VLAN trunk connected to a switch employing IEEE 802.1Q VLAN trunking
Stacking Link	A link that directly connects two VC ports from the same VC domain VC Virtual Connect: Broad term used to reference all the Virtual Connect components as a whole – Ethernet & Fiber Channel modules and Virtual Connect Manager.
Virtual Connect Domain	All VC Fiber Channel modules and all stacked VC-Enet modules within the same enclosure and under the control of the same Virtual Connect Manager.
Virtual Connect Network (vNet)	A logical grouping of VC ports (downlinks or downlinks & uplinks) that VC Downlink comprise a single layer 2 network or broadcast domain
VC-Enet	A Virtual Connect Ethernet module
VC-FC	A Virtual Connect Fiber Channel module
VCM	Virtual Connect Manager: The user interface, web or CLI, used to manage a Virtual Connect domain
VC Downlink	Non-visible ports that are directly connected to server NIC ports through the enclosure midplane.
VCEM	Virtual Connect Enterprise Manager: A separate software product that extends management to as many as 100 VC domains from a single console.
VC Uplink	Visible ports on the VC-Enet module faceplate that provide external connectivity for the enclosure.
VLAN Trunk	A single physical port or a single port channel with VLAN tagging enabled. Used to provide connectivity to one or more VLANs over the same logical path.
XFP	A hot-pluggable modular 10 Gbit port. Pluggable modules allow for electrical or optical connectivity at 10 Gbit speeds

For additional details – Refer to:

HP VC

HP BladeSystem Architecture Overview

The HP BladeSystem c-Class system includes an enclosure, server blades, storage blades, PCI expansion blade, interconnect modules (switches, Virtual Connect, and pass-thru modules), a NonStop signal midplane that connects blades to the interconnect modules, a shared power backplane, power supplies, fans, front LCD screen, and enclosure management controllers

(Onboard Administrator modules). The BladeSystem c-Class uses redundant and hot-pluggable components extensively to provide maximum uptime to the enclosure. Figure 1 shows the c7000 implementation of the architecture.

The HP BladeSystem architecture is available in two enclosure form factors: the c7000 and the c3000. The HP BladeSystem c7000 enclosure will accommodate up to 16 half-height server or other device blades, or up to eight full-height server blades, or a combination of the two blade form factors. The c7000 also provides eight rear bays for interconnect modules. The HP BladeSystem c3000 enclosure will accommodate up to 8 half-height server or other device blades, or up to 4 full- height server blades, or a combination of the two blade form factors. The c3000 also provides four rear bays for interconnect modules.

Optional mezzanine cards within the server blades provide network connectivity to the interconnect modules. The connections between server blades and the network fabric can be fully redundant. Customers install their choice of mezzanine cards in the mezzanine slots inside the blades, and interconnect modules in the interconnect bays at the rear of the enclosure to create their network fabric connectivity. Mezzanine card options include Ethernet, Fibre Channel, and Infiniband. However, Fibre Channel and Infiniband are outside the scope of this paper.

The enclosure houses either one or two Onboard Administrator modules. Onboard Administrator provides enclosure-base management and intelligence throughout the infrastructure to monitor power and thermal conditions, ensure hardware configurations are correct, and simplify network configuration. The Insight Display panel on the front of the enclosure simplifies configuration and maintenance. Customers have the option of installing a second Onboard Administrator module that acts as a completely redundant controller in an active-standby mode.

The c7000 enclosure can use either single-phase or three-phase power inputs and can hold up to six 2250 W power supplies. The power supplies connect to a passive power backplane that distributes the power to all the components in a shared manner.

To cool the enclosure, HP designed a fan known as the Active Cool fan. The c7000 enclosure can hold up to ten hot-pluggable Active Cool fans. The Active Cool fans are designed for high efficiency and performance to provide redundant cooling across the enclosure as well as providing ample capacity for future cooling needs.

Figure 1. Overview of c7000 Enclosure Components

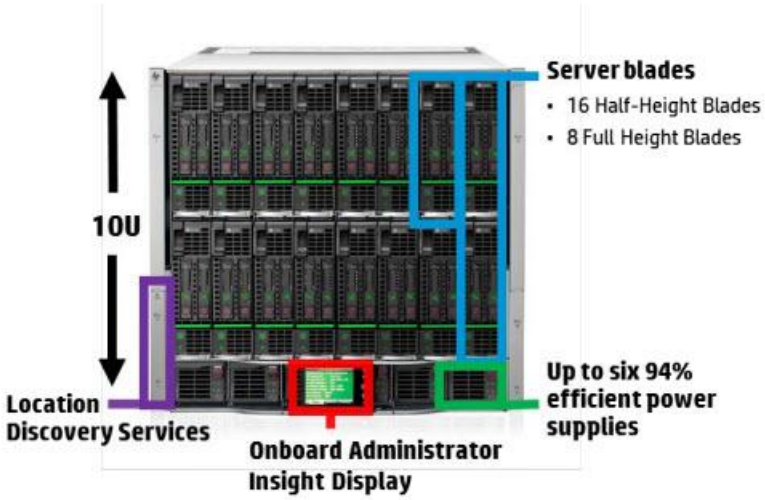

Figure 2. Rear view of c7000 Enclosure Components

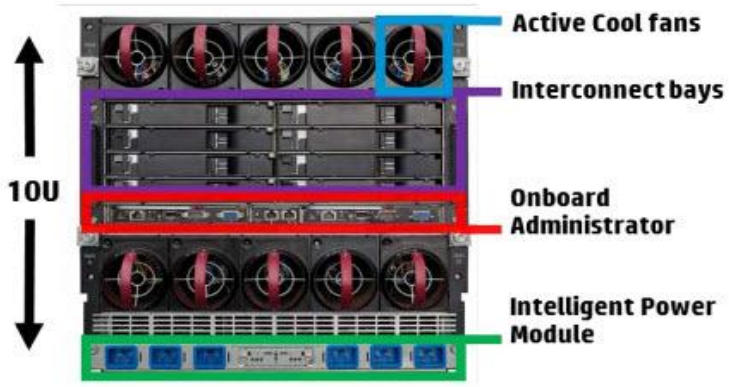

Figure 3. Midplane view c7000 Enclosure Components

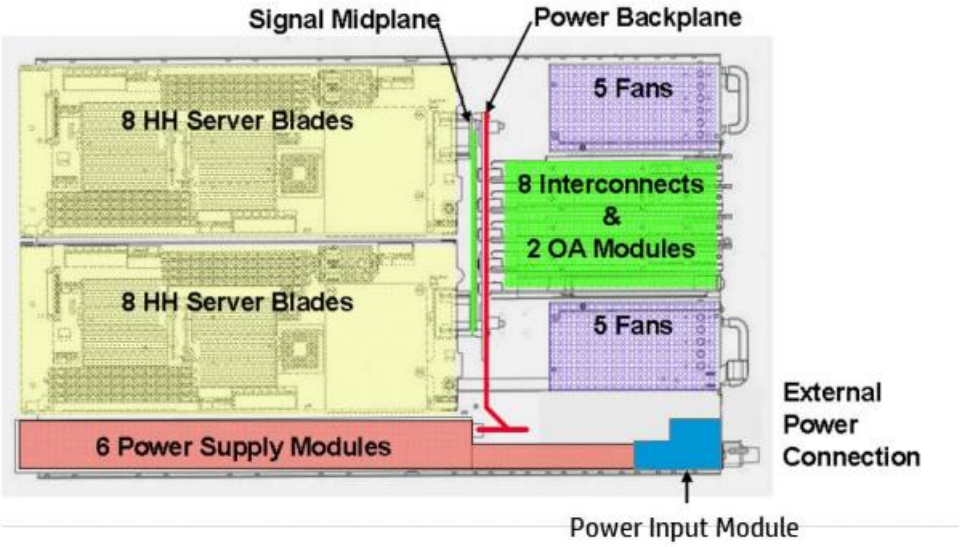

c-Class Cabling Layout

The c-Class server blades use various mezzanine cards to connect to different network fabrics through the interconnect bays at the rear of the enclosure. These fabrics include Ethernet, Fiber Channel, Infiniband, etc. The mezzanine cards are inserted in the mezzanine slots in the blade server. These mezzanine slots (M1, M2, & M3 in Figure 2 below) are basically PCI-e expansion slots equivalent to those found in non-bladed hardware such as tower and rack-mounted servers. The connections between the multiple types of mezzanine cards on the server blades are hard-wired through the NonStop signal midplane to specific interconnect bays at the rear of the enclosure (see Figure 2).

All c-Class blade servers have at least two onboard (LOM) NIC ports. These LOM ports, two for half- height servers (NIC 1 & NIC 2) and four for full-height servers (NIC 1 thru NIC4), are connected to interconnect bays 1 & 2 (see Figure 2). Since Ethernet LOM ports are the only type network fabric that are ever connected to interconnect bays 1 & 2, the only type of interconnect module that can be used in bays 1 & 2 are Ethernet modules (switches, pass-through, or Virtual Connect).

For half-height servers, mezzanine slot 1 (M1) provides up to two Ethernet ports connected to interconnect bays 3 & 4. Mezzanine slot 2 (M2) provides up to four Ethernet ports connected to interconnect bays 5 thru 8. The maximum configuration for a half-height server, if using all Ethernet, is a total of eight NICs where each NIC port is connected to one of the eight interconnect bays.

For full-height servers, mezzanine slot 1 (M1) provides up to four Ethernet ports connected to interconnect bays 3 & 4. Mezzanine slot 2 (M2) provides up to four Ethernet ports connected to interconnect bays 5 thru 8. Mezzanine slot 3 (M3) also provides up to four Ethernet ports which are also connected to interconnect bays 5 thru 8. The maximum configuration for a full-height server, if using all Ethernet, is a total of 16 NICs, where two NIC ports are connected to each of the eight interconnect bays.

Mezzanine slots may be populated with I/O technologies other than Ethernet (such as Fibre Channel or InfiniBand), but this paper focuses primarily on Ethernet.

C7000 Backplane Wiring

Figure 4. Overview of c7000 midplane wiring – Half height modules

Half Height Modules

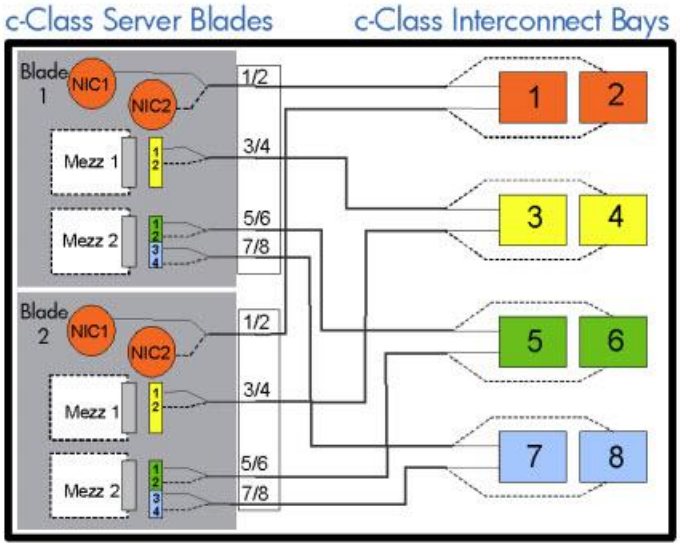

Figure 5. Overview of c7000 midplane wiring – Full height modules

Full Height

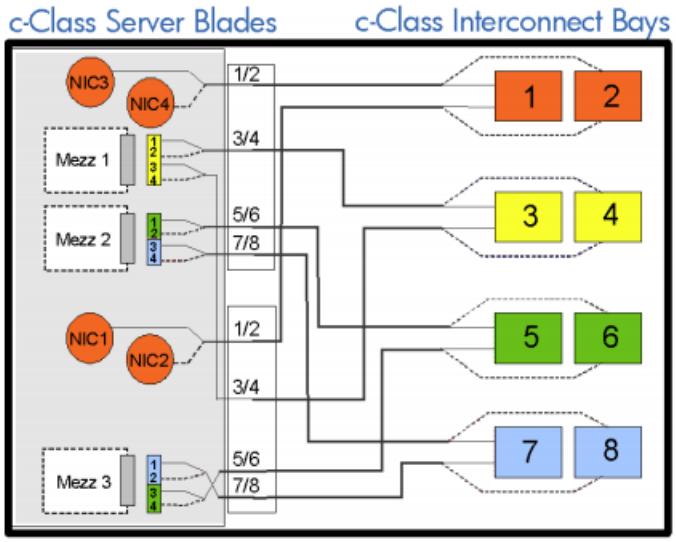

FlexNIC capabilities

Flex-10 and FlexFabric adapters allow you to partition a 10Gb link into several smaller bandwidth FlexNICs. Virtual machine applications often require increased network connections per server, increasing network complexity while reducing the number of server resources. Virtual Connect addresses this issue by letting you divide a 10Gb network connection into four independent FlexNIC server connections (Figure 4). A FlexNIC is a physical PCIe function (PF) that appears to the system ROM, OS, or hypervisor as a discrete physical NIC with its own driver instance. It is not a virtual NIC contained in a software layer.

You can configure bandwidth on each FlexNIC from 100 Mb up to 10Gb, and dynamically adjust the bandwidth in 100 Mb increments without requiring a server reboot. You can provide just the right amount of bandwidth based on application needs. You no longer need to over-provision or under-provision bandwidth. By virtue of Bandwidth optimization by setting Min and Max values for individual FlexNICs, Virtual Connect allocates unused bandwidth from FlexNICs to those FlexNICs whose bandwidth demands exceeds minimum. Min assures guaranteed bandwidth all the time and Max is best effort depending upon available bandwidth in other FlexNICs

Virtual Connect tells the network adapter how to configure each of the four physical functions. Then the adapter defines each of those physical functions, provisions them into the OS as individual NICs, and allocates the appropriate bandwidth. We work with each NIC vendor to ensure that they meet our Virtual Connect requirements for splitting the PCIe function and allocating bandwidth to each physical function.

Figure 6: Flex-10 adapters allow administrators to partition bandwidth based on application requirements.

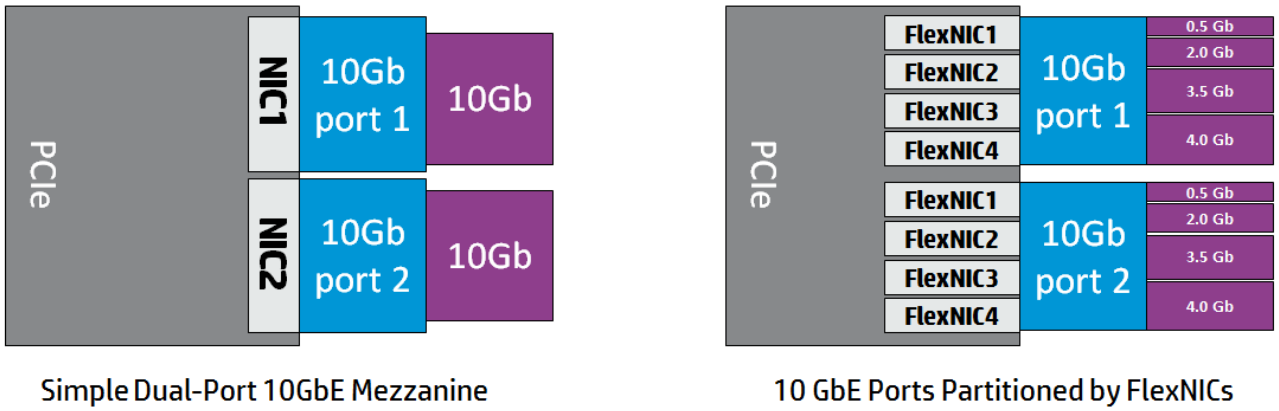

Figure 7: Flex-10/10D and Flex Fabric adapters allow administrators to partition bandwidth based on application requirements – including the ability to support FCOE SAN traffic.

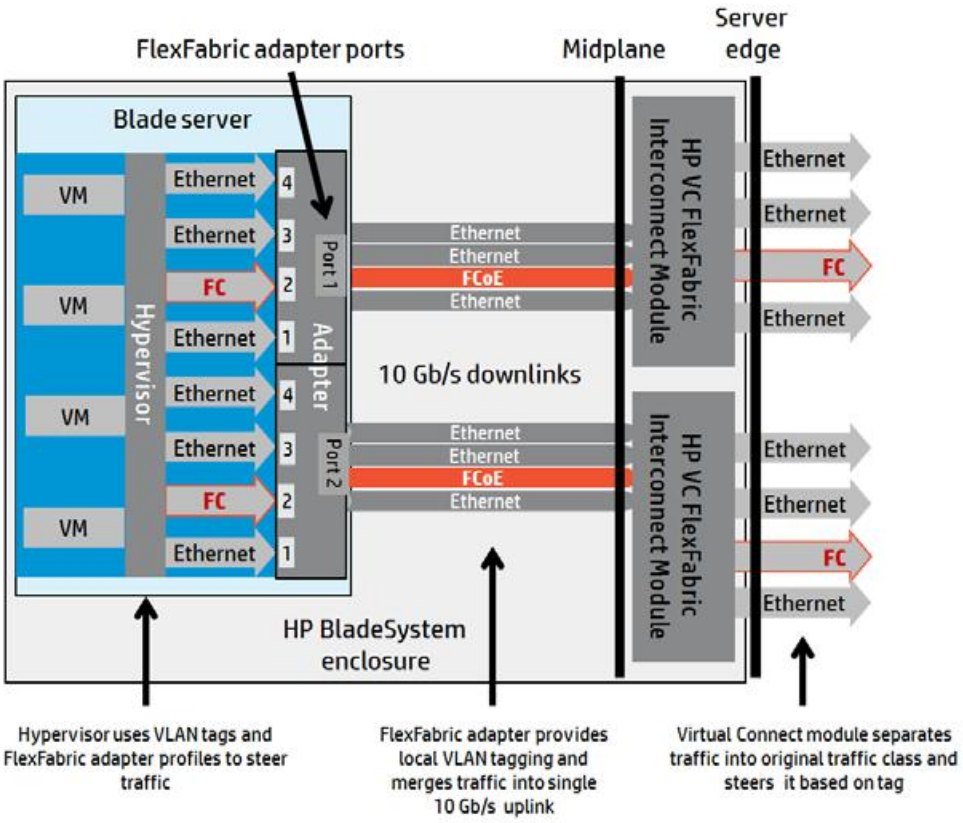

c-Class Ethernet Interconnect Options

The BladeSystem c7000 Enclosure offers a variety of interconnect options, including pass-thru modules, Ethernet and Fibre Channel switches, Virtual Connect modules, and high-bandwidth fabrics such as InfiniBand. The HP website (www.hp.com/go/bladesystem/interconnects) contains the most up-to-date information about the available c-Class interconnect modules.

BladeSystem interconnect modules incorporate a variety of networking standards:

- Ethernet
- Fibre Channel
- Fibre Channel over Ethernet (FCoE)
- InfiniBand
- iSCSI
- Serial Attached SCSI (SAS)

Ethernet Switches from Cisco HP Networking and Mellanox offer customers a traditional approach to administering the network. The primary value in blade switches is cable consolidation through high-speed uplinks and the shared blade power and cooling infrastructure.

Available Blade "Switch" Options:

- *Cisco Catalyst Blade Switch 3020*
- *Cisco Catalyst Blade Switch 3120*
- *Cisco Fabric Extender B22HP*
- *HP Ethernet Blade Switch 6125XLG*
- *HP Ethernet Blade Switch 6125G*

- *HP Ethernet Blade Switch 6120G*
- *HP Ethernet Blade Switch 6120XG*
- *HP GbE2c Layer 2/3 Ethernet Blade Switch*
- *HP 1:10Gb Ethernet BL-c Switch*
- *Mellanox Ethernet Switch SX1018HP*

Ethernet pass-thru modules are also available when direct one-to-one connections between servers and the LAN is required. HP Ethernet Pass-Thru Modules provide 16-port, transparent, 1:1 port connectivity between the server and an external Ethernet switch.

Finally, Virtual Connect is also an option for customers for use in place of c-Class Ethernet switches or Ethernet pass-thru modules. Further discussions about the feature sets of Virtual Connect will follow in the subsequent pages.

HP Virtual Connect Technology

Overview

Virtual Connect is a set of interconnect modules and embedded software for HP BladeSystem c-Class enclosures that simplifies the setup and administration of server connections. The HP Ethernet Virtual Connect portfolio includes the HP 10/10D and FlexFabric Modules for c-Class BladeSystem, The Fibre Channel portfolio is filled out by the HP 20 and 24 port 8Gb Virtual Connect Fibre Channel Modules. All modules are managed with the embedded HP Virtual Connect Manager, the optional Virtual Connect Enterprise Manager (VCEM) or the new optional Oneview Manager. Virtual Connect implements server edge virtualization so that server administrators can upgrade, replace, or move server blades within their enclosures without changes being visible to the external LAN and SAN environments.

The Virtual Connect Manager is embedded on the HP 10/10D / FlexFabric Virtual Connect Ethernet Modules. Virtual Connect can be accessed by Administrators through web links provided by the Onboard Administrator, through direct connection to the embedded Virtual Connect Manager web server, through a serial connection through the OA for CLI access, or through an SSH session for CLI access to the Virtual Connect Manager.

A Virtual Connect Domain, configured using the Virtual Connect Manager, includes a single HP c- Class BladeSystem enclosure, all installed Virtual Connect modules and a total of 16 server bays that can be populated with a mixture of server and or storage blades. Within the VC domain, any server blade can be configured to access any LAN or SAN connected to a VC module, and a server blade can be used as a spare for any server blade within the same VC domain.

By stacking (cabling) the Ethernet modules within the VC domain, every server blade in the domain can be configured to access any external network connection. Fibre Channel modules (VC-FC) within different I/O bays are each connected directly to individual FC SAN fabrics. With this configuration, the Virtual Connect Manager can deploy and move a server profile to any server bay within the Virtual Connect domain without the need to change external LAN or SAN configurations.

The Virtual Connect Enterprise Manager is a separate software product that extends management to as many as 250 VC domains from a single console, provides a single pool of managed MAC Addresses and World Wide Names, enables movement of Virtual Connect server profiles from one

VC Domain to another, and enables the automatic failover of Virtual Connect server profiles from one server bay to another. Oneview exponentially expands the capabilities of connected enclosures. This paper references the embedded Virtual Connect Manager, and not Virtual Connect Enterprise Manager and or Oneview, in its explanations and examples.

More information on VCEM is available at:

VCEM

More information on Oneview is available at:

Oneview

Virtual Connect Components

There are three key components that make up the Virtual Connect infrastructure. Two components are hardware and one component is software. The three key components are Virtual Connect Ethernet modules, Virtual Connect Fibre Channel modules, Virtual Connect Manager (VCM), and the optional Virtual Connect Enterprise Manager (VCEM) and or the new Oneview application.

Virtual Connect Ethernet Modules:

Virtual Connect Flex-10/10D Module

It is important to note how the external uplink ports on the Flex-10/10D module are configured. The graphic below outlines the type and speed each port can be configured as:

- Ports X1 – X10; Can be configured as 1Gb or 10Gb Ethernet or FCoE (ALL external ports can be used, no sharing of these ports with internal stacking, as with previous modules)
- Ports X11-X14; Internal cross connections for horizontal stacking and are NOT shared with any external connections
- Uplink Ports X1-X10 support 0.5–15m length DAC as stacking or uplink. If greater lengths are required, fibre optic cables would be required

Figure 8: Flex 10/10D Module

Virtual Connect 10/10D Module

For latest specifications:

Flex 10/10D

Virtual Connect FlexFabric Module

It is important to note how the external uplink ports on the Flexfabric module are configured. The graphic below outlines the type and speed each port can be configured as:

- Ports X1 – X8; Can be configured as 1Gb or 10Gb Ethernet or FCoE
- Ports X7 and or X8; are used for cross connections for horizontal stacking. Only one port is needed – two if you want redundancy.
- Uplink Ports X1-X8 support 0.5–15m length DAC as stacking or uplink. If greater lengths are required, fibre optic cables would be required

Figure 9: FlexFabric Module

8 Port Flexfabric Module

For additional detailed information:

Flexfabric

Virtual Connect Fibre Channel Modules

Virtual Connect 24 Port Modules

- Virtual connect Fiber Channel modules come in a 24 port and port configuration.
- Both modules support 4 and 8 Gig Fibre Channel SFP's.
- VCFC modules are managed by ether Flex10/10D or Flex Fabric modules running the active VCM

Figure 10: Virtual Connect 24 Port Fibre Channel Module

24 Port Virtual Connect Fibre Channel module For more detailed information: 24 Port VCFC

Virtual Connect 20 Port Module

Figure 11: Virtual Connect 20 Port Module

20 Port Virtual Connect Fibre Channel module For more detailed information: 20 Port VCFC

Virtual Connect Manager

The purpose of the Virtual Connect Manager (VCM) is to function as the single point of administration for the Virtual Connect Domain. This means that all aspects of configuring, managing, and monitoring the Virtual Connect Domain and all VC Ethernet and Fiber Channel modules is provided by VCM.

VCM runs embedded on either the HP 10/10D or the HP Flex Fabric Ethernet modules when installed in interconnect any side by side interconnect bay. VCM utilizes its own CPU, memory, and flash resources onboard the VC Ethernet modules.

VCM runs in High Availability mode when Virtual Connect Ethernet modules are installed in both side by side interconnect bay. By default, the Virtual Connect Ethernet module in odd interconnect bay will be the active VCM and the Virtual Connect Ethernet module in even interconnect bay will be the standby VCM. The Virtual Connect Domain configuration is check pointed from the active VCM to the standby whenever configuration changes are made. In this way, if the Virtual Connect Ethernet module functioning as the active VCM fails, the standby Ethernet module can assume the role of active VCM without loss of Virtual Connect Domain configuration settings.

Note:
Even though VCM is only active on a single Virtual Connect Ethernet module at one time, this has nothing to do with the Ethernet connectivity functionality provided by the Virtual Connect Ethernet modules. All Virtual Connect modules in the domain are functionally active at all times and every port on every Virtual Connect module can be configured to actively carry network traffic at the same time.

Virtual Connect Manager User Interfaces

Virtual Connect Manager (VCM) provides three user interfaces for managing and monitoring the Virtual Connect Domain – Web, SSH CLI, and SNMP.

The Web and SSH CLI interfaces provide nearly an identical set of features used to manage the Virtual Connect Domain. For more information on the specifics of these interfaces, please consult the Virtual Connect Users Guide: (www.hp.com/go/bladesystem/documentation).

VC's SNMP implementation allows the administrator to monitor the Virtual Connect domain and associated modules. The SNMP configuration is controlled by the Virtual Connect Manager and applies to all modules in the VC domain. The SNMP agent software module residing on a primary module provides access to managed information using SNMP protocol. The managed information is presented as a hierarchical database called a MIB. Each element of the managed information is identified by a unique identifier called Object ID. Basic SNMP support is provided for VC-Enet modules.

VC MIB Support

VC modules support the applicable groups of the following SNMP MIBs:

Figure 12:

MIB	VC-Enet	VC-FC
RFC 2863 IF-MIB	X	—
RFC 4188 Bridge-MIB	X	—
RFC 3418 SNMP v2 MIB	X	X
Compaq System Info MIB	X	X
Compaq Host MIB	X	X
Compaq Rack MIB	—	X*
RFC 1213 Network Mgmt	X	—
RFC 4293 IP-MIB	X	—
Fibre Alliance MIB (FC Mgmt Integ)	—	X
RFC 2837 Fabric Element MIB	—	X
VC Module MIB (VCM-MIB)	X	—
VC Domain MIB (VCD-MIB)	X	—

Virtual Connect Port Descriptions

The three types of ports are VC downlinks, VC uplinks, and internal cross-connects:

VC Downlinks

Ports that are directly connected to server NIC ports through the enclosure midplane. Only role is to provide connectivity to directly connected blade server NICs

VC Uplinks

Ports on the VC-Enet module faceplate that provide external connectivity for the VC domain. Roles include stacking link, network analyzer port (see Port Monitoring), normal mode (providing external connectivity for one or more vNets or Shared Uplink Sets)

Internal cross-connects

A non-visible port that interconnects two horizontally adjacent VC-Enet modules Only role is to function as a stacking link.

Overview of Virtual Connect Ethernet Connectivity

While Virtual Connect provides many features, its primary purpose is to provide internal and external network connectivity for blade servers in a c-Class enclosure. For Ethernet, Virtual Connect accomplishes this purpose by bridging together, at layer 2, VC downlink ports (server NICs) with other VC downlink ports and with VC uplink ports connected to an external switched network. Virtual Connect provides the administrator with control over which uplinks and which downlinks are assigned to the same layer 2 network through the use of Virtual Connect Networks (vNets or Shared Uplink Sets). This mapping technology prevents or manages data traffic

to the various ports – IE ports that are not mapped cannot communicate.

A VC Administrator defines vNets and, optionally, assigns VC uplinks to the vNets to provide external network connectivity. Once the vNets are defined, they are available for the Server Administrator to assign to server NIC ports. When a Server Admin moves a server profile from one server bay to another, the vNet assignments and, if used, managed MAC addresses (see section entitled "VC Managed MAC Addresses) are moved with the profile. This allows the server administrator to easily move the "Ethernet connectivity profile" for a server from one bay to another without requiring the assistance of the Network Administrator.

A vNet can be used to bridge together multiple VC downlinks for an internal-to-the-enclosure-only network or a vNet can be used to bridge together one or more VC downlinks with one or more VC uplinks to provide external network connectivity for the blade servers in the enclosure. A vNet cannot be used to bridge together multiple VC uplinks to provide connectivity between two external devices or to provide connectivity between two external networks. Virtual Connect is not, and can not be configured as, a transit device. In other words, VC can only provide internal connectivity between blade servers and\or VC can provide external connectivity to blade servers but Virtual Connect cannot provide connectivity between two external devices.
The basic definition of a Virtual Connect Network, or vNet, is a logical grouping of VC ports (downlinks or downlinks & uplinks) that comprise a single layer 2 network or broadcast domain. A vNet is analogous to a vSwitch, as used by VMware technology (see "Comparing VC and VMware Networking Technology" for an in-depth comparison). A vNet does not always represent a one-to- one correlation with a VLAN, as can be seen in the next section entitled "Virtual Connect Ethernet Connectivity Designs".

Server NIC ports assigned to the same vNet can communicate directly (at layer 2) without having to leave the enclosure. In the figure below, Server Blade 1 and Server Blade 2 can communicate with each other within the VC Domain. Also, server NIC ports in different vNets are isolated (at layer 2) by Virtual Connect from other server NIC ports. Server Blade 16 is isolated from Server Blades 1 thru 3.

Figure 13. Sample VC Domain Showing Various Uses of vNets

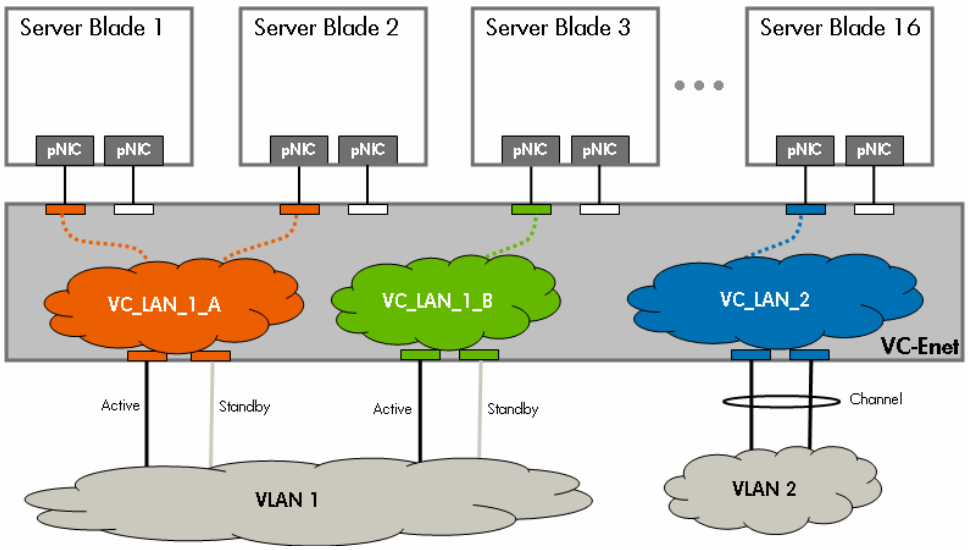

(see Appendix A for a description of the elements in the above diagram)

Note:
Layer 2 connectivity is defined as any two devices that can communicate with each other by directly exchanging Ethernet frames, carrying any protocol, without traversing a router or layer 3 switch. A layer 2 network could also be defined as a "broadcast domain". For example, when a frame is broadcast within a vNet, only ports assigned to the vNet will receive the broadcast frame (unless an external device bridges multiple vNets together).

VC Uplink Fault Tolerance

Virtual Connect can be configured to provide both fault tolerance and load balancing for Virtual Connect Networks and the associated server NIC ports. An Administrator can choose whether a vNet operates in fault tolerance only mode by setting the vNet's connection mode to "Failover" or the Administrator can choose fault tolerance plus load balancing by setting the vNet's connection mode to "Auto".

When an administrator assigns multiple VC uplinks to the same vNet, VC's default behavior (connection mode 'auto') for a vNet (or Shared Uplink Set) is to attempt to negotiate a port channel (EtherChannel) using 802.3ad Link Aggregation Control Protocol (LACP). If LACP negotiation fails, the vNet operates in fault tolerance mode only. Similar to the operation of NIC Teaming, Virtual Connect will utilize one VC uplink port as the active port and all other VC uplink ports will be in standby (blocking) mode. This behavior is part of the loop prevent mechanism of VC (see section entitled "Virtual Connect's Loop Prevention Technology"). See figure below as an example.

Figure 14. A vNet Configured for Fault Tolerance only

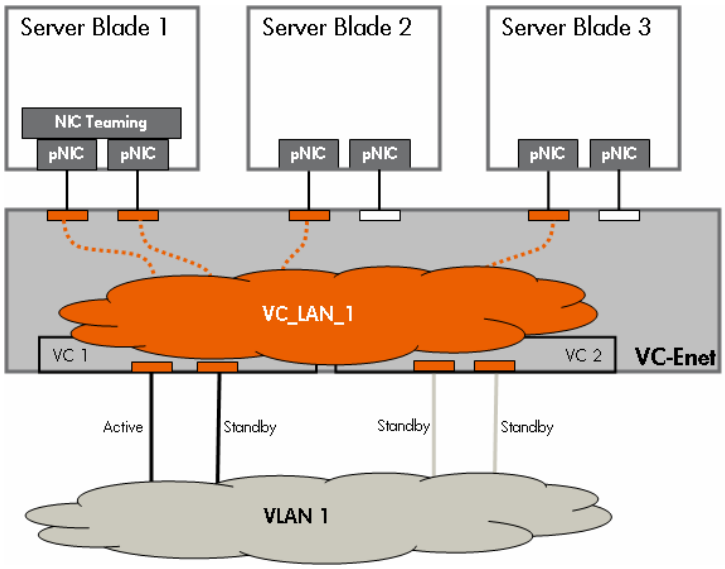

(see Appendix A for a description of the elements in the above diagram)

An administrator can choose to manually configure a vNet (or Shared Uplink Set) for **only** fault tolerance by setting the vNet's connection mode to 'failover'. When a vNet's connection mode is set to 'failover', the LACP protocol is disabled for the vNet and the administrator can change the VC uplink "Port Role" to "Primary" or "Secondary". This port role setting allows the administrator to manually choose which VC uplink should be the preferred active uplink.

Note:
For a sample configuration showing VC uplinks connected to a Cisco switch, see the section below entitled "Sample Virtual Connect Ethernet and Cisco Configurations"

A vNet (or Shared Uplink Set) will failover from one uplink to another whenever the active uplink loses link. As soon as VC detects link loss on the active uplink, VC will choose a new active uplink from the group of available standby uplinks also assigned to the same vNet. This process of choosing a new active uplink typically occurs in under five seconds. If the upstream Cisco switch has been properly configured on all the uplinks associated with the vNet, the server connectivity should be restored in under five seconds. If the previous active uplink has it's link restored, VC will automatically failback to it. As long as the external Cisco switch port has been properly configured by enabling PortFast, connectivity to servers should be restored in under five seconds.

Virtual Connect uses a mechanism called "Fast MAC Cache Failover" to proactively update the CAM tables on the upstream Cisco switch ports. Whenever a vNet fails over or fails back from one uplink to another, this mechanism will transmit a single frame on the new active uplink for every server MAC address that is active within the VC domain. Fast MAC Cache Failover enables VC to immediately update the external Cisco infrastructure with the new location of the active VC uplink for a particular vNet. Fast MAC Cache Failover is very similar to the CAM table update process used by Link-State Tracking on Cisco switches.

VC Uplink Load Balancing

Virtual Connect provides load balancing (and fault tolerance) across multiple VC uplinks on the same physical VC-Enet module by means of IEEE 802.3ad port trunking or port channeling (EtherChannel) using the Link Aggregation Control Protocol (LACP). The vNet (or Shared Uplink Set) must have its "connection mode" set to "auto" (default) and not in mode "failover".

If the external switch is a Cisco switch, the ports in the EtherChannel must be set to either "mode active" or "mode passive". Both of these modes on the Cisco switch enable the use of 802.3ad LACP (the only port channeling protocol VC supports). VC does not support port channeling with Cisco switch ports when set to "mode on", "mode desirable" or "mode auto".

Note:
An EtherChannel can only be formed between VC and a Cisco switch if the VC vNet is set to connection mode "Auto" and the Cisco switch's channel-group is set to either "mode active" or

"mode passive".

There are four types of scenarios for port trunking\channeling with Virtual Connect Uplink Ports:

Same VC Module, Same External Switch, Same Port Channel Group
When VC uplinks from the same physical VC-Enet module are assigned to the same vNet (or Shared Uplink Set) and are connected to ports on the same external Cisco switch that are assigned to the same port channel group and have the LACP protocol enabled, then VC will automatically form a single port channel.

Same VC Module, Same External Switch, Different Port Channel Groups
When VC uplinks from the same physical VC-Enet module are assigned to the same vNet (or Shared Uplink Set) and are connected to ports on the same external Cisco switch that are assigned to different port channel groups and have the LACP protocol enabled, then VC may automatically form more than one port channel.

Same VC Module, Different External Switches
When VC uplinks from the same physical VC-Enet module are assigned to the same vNet (or Shared Uplink Set) and are connected to ports on different external Cisco switches that are assigned to port channel groups and have the LACP protocol enabled, then VC may automatically form more than one port channel.

Different VC Module, Same or Different External Switches
When VC uplinks from different physical VC-Enet modules are assigned to the same vNet (or Shared Uplink Set) and are connected to ports on the same (or different) external Cisco switch that are assigned to a port channel group and have the LACP protocol enabled, then VC may automatically form more than one port channel.

As described above, VC's default behavior (connection mode 'auto') for a vNet is to attempt to negotiate a port channel (EtherChannel) using 802.3ad Link Aggregation Control Protocol (LACP). If LACP negotiation is successful for one or more sets of VC uplink ports, a port channel is formed between the VC Ethernet module and the external switch. A vNet may form more than one port channel with an external switch or switches. Each port channel behaves as a single logical path. Like in the previous section describing simple fault tolerance, a vNet can only have one active logical path at any given time. As such, a vNet with multiple port channels will only use one port channel as the active logical path and all other port channels in the same vNet will be in standby mode. See the figure below as an example.

Figure 15. A vNet Utilizing Port Channeling for Load Balancing and Fault Tolerance

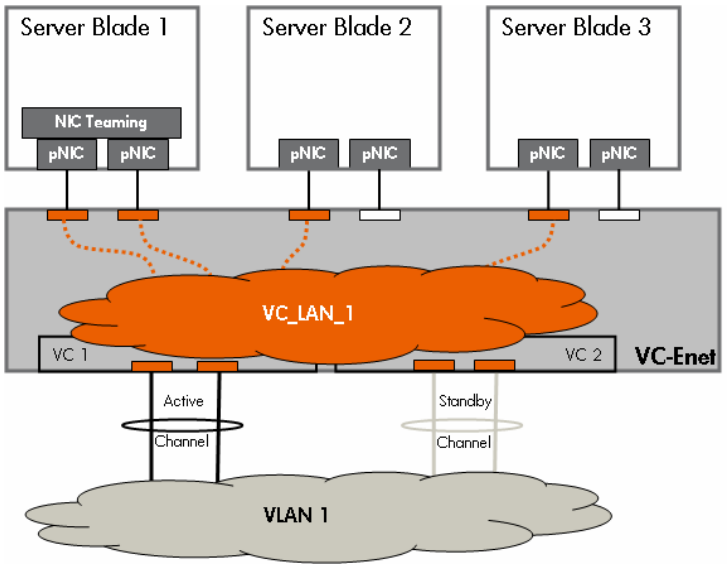

(see Appendix A for a description of the elements in the above diagram)

Note:
Port channels can only form when VC uplink ports on the same physical VC Ethernet module are connected to the same external switch. VC does not support port channels that span across different VC Ethernet modules.

Port Channeling (802.3ad) Load Balancing Algorithm

VC-Enet's implementation of port channeling, or EtherChannel, uses a load balancing algorithm for the purpose of distributing frames across the physical ports that make up the port channel. The biggest concern in distributing frames across multiple physical ports is "frame ordering". For any particular conversation between two network devices (e.g. FTP transfer, telnet session, etc.), it is optimal that the network infrastructure deliver the frames in the order in which the transmitter transmitted them. This minimizes frame reordering on the receivers end. Also, while TCP provides header information for putting frames back into the correct order, other protocols, such as UDP, do not. Therefore, frame ordering is critical. This means that any load balancing algorithm used by port channeling must load balance frames but also maintain frame ordering. This is accomplished by using an algorithm that makes the same load balancing decision for frames from the same conversation.

Typical EtherChannel Load Balancing Algorithms (from worst to best):

- **Source MAC address**

  Identifies all conversations coming from the same MAC address and load balances them all down a single link in the port channel

- **Destination MAC address**

  Identifies all conversations destined for the same MAC address and load balances them all down a

single link in the port channel

- **Source IP address**

  Identifies all conversations coming from the same IP address and load balances them all down a single link in the port channel

- **Destination IP address**

  Identifies all conversations destined for the same IP address and load balances them all down a single link in the port channel

VC

- **Source and Destination MAC address**
  Identifies all conversations between two MAC addresses and load balances them all down a single link in the port channel
- **Source and Destination IP address**
  Identifies all conversations between two IP addresses and load balances them all down a single link in the port channel
- **Source and Destination IP address plus TCP or UDP ports (socket)**
  Identifies specific conversations between two IP addresses and may load balance them down different links in the port channel

Virtual Connect utilizes one of the last three algorithms depending on the type of frame being load balanced.

For unicast frames, VC load balances using the socket (best method above) if it has a TCP or UDP header, by using the source and destination IP address if it is an IP frame without a TCP or UDP header, or by using the source and destination MAC address if it is not an IP frame.

Broadcast, multicast, and unknown unicast (destination lookup failure) frames are handled slightly differently than unicast frames. They are load balanced so VC does not always send these frames down the first port in the channel. The algorithm used to determine the port in the channel is based on the source and destination MAC addresses, the source port channel number (if source is a port channel) and/or module ID plus port number (if source is a regular port and not a port channel).

Since VC automatically load balances traffic across the ports in a port channel using the most conversation-specific information available (TCP/UPD, then IP, then MAC), VC does not provide a user configurable setting for this feature. Also, the load balancing algorithm does not have to match on both sides of the same port channel. In other words, VC can use one algorithm while the directly connected Cisco switch uses a different algorithm. Each side can implement its own algorithm without affecting connectivity. As a result, while HP advises that the directly connect Cisco switch utilize a load balancing algorithm based on source-destination IP plus TCP\UDP port, it is not a requirement. If the external Cisco switch uses an algorithm based on source-destination MAC or source-destination IP address, it will also work with Virtual Connect.

Note:
For a sample configuration showing VC uplinks connected to a Cisco switch configured for EtherChannel, see the section below entitled “Sample Virtual Connect Ethernet and Cisco Configurations”

VC Uplinks and VLAN Trunking

Virtual Connect supports VLAN tagging (trunking) on VC uplinks using IEEE 802.1Q and can be configured to support VLAN tagging on blade server NIC ports.

VC Uplinks can operate in one of three modes:

- **VLAN trunking mode**
  - Shared Uplink Set (SUS) plus tagging on external switch port
  - When a VC uplink is a member of a SUS and is connected to a VLAN trunk port on an external switch, it operates as a VLAN trunk port.
  - Multiple vNets, multiple VLANs
  - No tagging on server NIC ports
  - Example: VC Uplink 3 in figure below
- **VLAN tunneling mode**
  - Single vNet plus tagging on external switch port
  - When a VC uplink is a member of a single vNet and is connected to a VLAN trunk port on an external switch, it operates in a mode that is equivalent to 'dot1qtunnel' mode on a Cisco switch.
  - Single vNet, multiple VLANs
  - Tagging on server NIC ports
  - Example: VC Uplink 2 in figure below
- **Access mode**
  - Single vNet plus no tagging on external switch port
  - When a VC uplink is a member of a single vNet and is connected to an external switch port in access mode, it operates in a mode that is equivalent to 'access' mode on a Cisco switch.
  - Single vNet, single VLAN
  - No tagging on server NIC Ports
  - Example: VC Uplink 1 in figure below

When VC is configured to support VLAN tagging on a VC uplink (VLAN trunking mode), this means that VC is interpreting (i.e. removing and inserting tags, classifying the frame's vNet membership based on tag info) the tag and is a termination point for the tag (i.e. the servers do not see the VLAN tag). This is accomplished on VC by using "Shared Uplink Sets", or SUS for short. A SUS is one or more VC uplinks that are configured as VLAN trunks (tagging) and correlate external VLANs to specific vNets within the VC Domain. The Administrator creates the SUS, adds one or more VC uplinks, defines which VLAN tags are carried on the VC uplinks and then provides a vNet name (ASCII) for each of the VLANs defined. This means that the uplink ports within the SUS provide external connectivity for multiple vNets. In other words, one or more vNets share a set of uplinks for external network connectivity. This is why the term "Shared Uplink Sets" is used. See VC Uplink 3 in the figure below.

When VC uplink ports are assigned to a single vNet and connected to an external switch port in VLAN trunking mode, the VC uplink and vNet operate in VLAN\CoS tunneling mode or dot1qtunnel mode by

default. In other words, the vNet keeps all frames within the same layer 2 domain (vNet), however, VC allows the frames to carry different VLAN tags from the external network all the way to the server NIC ports and vice versa. See VC Uplink 2 in the figure below. When VC uplink ports are assigned to a single vNet and connected to an external switch port in access mode, the VC uplink and vNet operate in access mode only carries untagged frames for a single VLAN (the access VLAN on the switch). See VC Uplink 1 in the figure below. In other words, depending on the configuration of the external switch port, a single vNet could represent a single VLAN or a single vNet could represent multiple VLANs.

Figure 16. VLAN Tagging Examples

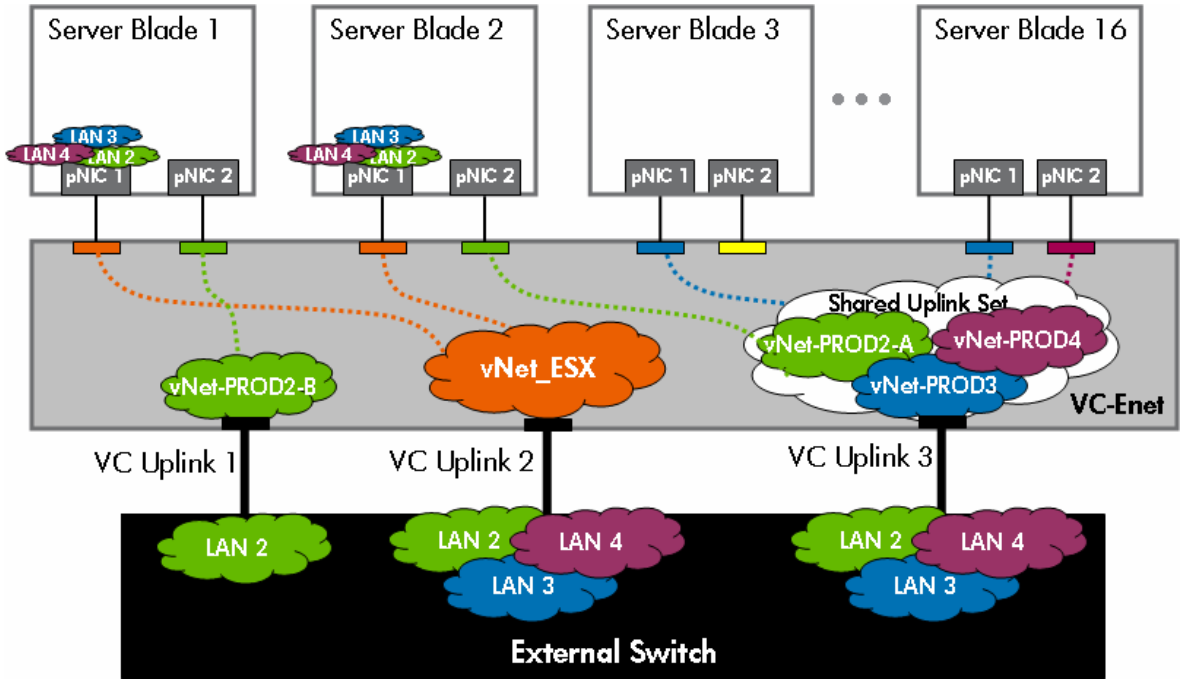

(see Appendix A for a description of the elements in the above diagram)

Explanation of the Figure Above:

VC Uplink 3 is assigned to a Shared Uplink Set (SUS) and is tagged with VLANs 2 thru 4.
The external switch port is a trunk port carrying VLANs 2 thru 4
VLAN 2 is represented by vNet-PROD2-A, VLAN 3 is represented by vNet-PROD3, and VLAN 4 is represented by vNet-PROD4.

VC Uplink 2 is assigned to vNet_ESX and is operating in 802.1Q tunneling mode.
The external switch port is a trunk port carrying VLANs 2 thru 4.
vNet_ESX is a single vNet that represents VLANs 2 thru 4.

VC Uplink 1 is assigned to vNet-PROD2-B and is operating in access mode.
The external switch port is an access port assigned to VLAN 2.
vNet-PROD2-B represents a vNet operating in 'acccess mode' for external VLAN 2.

- NIC 1 on both Server 1 and Server 2 have VLAN tagging configured on the server. On

both servers, NIC 1 represents VLANs 2 thru 4.

- All other NICs are not using VLAN tagging. All of these NICs are connected to a single VLAN (see color for VLAN association). NIC 1 on Server 1 can talk directly to NIC 1 on Server 2 without leaving the VC Domain. In order for NIC 2 on Server 1 to talk to NIC 2 on Server 2, the frames must first exit the VC Domain via VC uplink 1, transit the external switch on VLAN 2, and then re-enter the VC Domain on VC uplink 3.

Shared Uplink Sets manage uplink redundancy the same as individual vNets do. In other words, when multiple uplinks are assigned to a Shared Uplink Set, the uplinks can operate in failover-only mode or they can operate in port channeling (EtherChannel) mode. Also, all VLANs and associated vNets within a single Shared Uplink Set use the same active uplink or same active port channel.

Note:
For a sample configuration showing VC uplinks connected to a Cisco switch configured for VLAN trunking, see the section below entitled "Sample Virtual Connect Ethernet and Cisco Configurations". The Cisco switch commands are the same for VC in VLAN Trunking mode or VC in VLAN Tunneling mode.

Unsupported Cisco VLAN Protocols

There are two proprietary Cisco VLAN protocols that Virtual Connect does not support – VTP and DTP.

VTP, or VLAN Trunking Protocol, is a Cisco proprietary protocol used to manage a single VLAN database across multiple switches within the same domain. Since VC is not a Cisco device, VC does not support VTP. As such, there is no possibility that an improperly configured VC Domain can cause harm to the VTP Domain. Since VC does not support VTP, an administrator is required to add every VLAN that needs to be trunked on a VC uplink port.

DTP, or Dynamic Trunking Protocol, is a Cisco proprietary protocol used for automating the VLAN trunking configuration between a pair of directly connected Cisco ports. With DTP enabled, one Cisco switch port can be configured for VLAN trunking and the partner switch port can automatically configure its VLAN trunking configuration to match. Since VC is not a Cisco device, VC does not support DTP. HP recommends that the Administrator disable DTP on the directly connected Cisco switch ports using the command "switchport nonegotiate".

Virtual Connect and NIC Teaming

Virtual Connect supports NIC Teaming (or NIC bonding) on server NIC ports. For Windows on x86, VC supports Network Fault Tolerance (NFT) and Transmit Load Balancing (TLB) but does not support Switch-assisted Load Balancing (SLB). For Windows on Integrity, VC supports Network Fault Tolerance (NFT), Transmit Load Balancing (TLB), and static Dual Channel with only two NIC ports in the team, but does not support Switch-assisted Load Balancing (SLB). For Linux, VC supports any NIC bonding type that does not require 802.3ad (static or dynamic using LACP) on the

server NIC ports.

Virtual Connect also supports a feature called "SmartLink" that is enabled on vNets used by servers with NIC Teaming\bonding enabled. The SmartLink feature will disable the VC downlinks (server NIC ports) whenever all the VC uplinks for the associated vNet are unplugged. Basically, the SmartLink feature propagates a link-down event of a vNet's uplinks to the server NICs in the team. This allows the NIC Teaming software to failover from one vNet (with all failed uplinks) to another vNet with functional uplinks. SmartLink is only used in NIC Teaming configurations when the NIC ports in the Team are connected to different vNets (see "Figure 12. VC Domain Showing Advanced Usage of vNets and EtherChannel" for an example). HP recommends that SmartLink only be enabled on a vNet if ALL server NIC ports assigned to the vNet are using NIC Teaming and are connected to at least one other vNet.

For additional information on NIC Teaming for Windows, see link in the "Additional Resources" section at the end of this paper.

Virtual Connect and Cisco Terminology Comparison

In addition to using many traditional networking terms, Virtual Connect also introduces some new networking terminology. Some new terms are needed to properly describe VC's functionality since VC represents a new way of interconnecting blade servers to the core network. To aide in quickly understanding these Virtual Connect networking terms, the following table compares VC and Cisco networking terms.

Figure 17: Virtual Connect and Cisco Port Trunking and VLAN Trunking Terminology Comparison

Industry Common	Cisco Catalyst	HP Virtual Connect	Diagram
VLAN Trunking (802.1Q)	Trunking (ISL* or 802.1Q)	Shared Uplink Set (802.1Q)	
Port Trunking (static, 802.3ad LACP)	EtherChannel (static, PAgP*, or 802.3ad LACP)	Auto (802.3ad LACP)	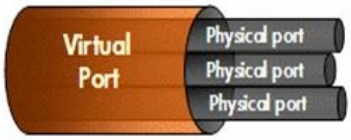
VLAN Trunking - Plus - Port Trunking	Trunking - Plus - EtherChannel (* Cisco proprietary)	Shared Uplink Set - Plus - Auto	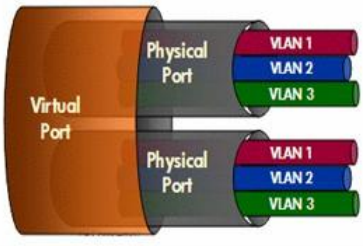

Sample Virtual Connect Ethernet and Cisco Nexus Configurations

The sample below is the most common configuration request that we receive.

Sample Configuration: VC Uplinks Connected to Cisco Nexus Access Ports

Additional Configuration details can be found at:

Dual Hop FCOE with Virtual Connect

HP Virtual Connect FlexFabric Cookbook

HP Virtual Connect 1 GB Cookbook

HP Virtual Connect for c-Class Blade System

Figure 17: Dual-Hop FCOE configuration

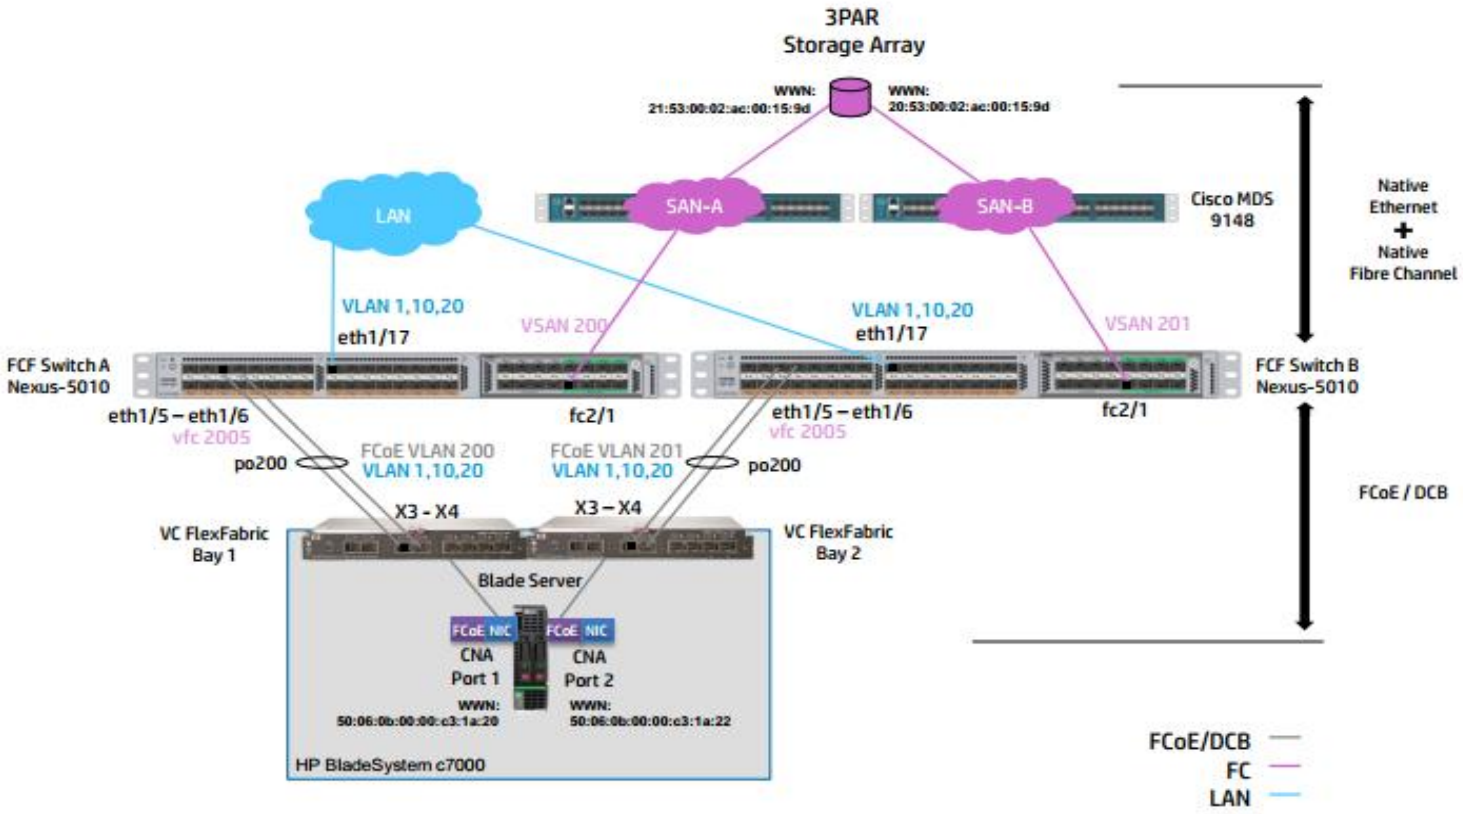

Figure 17 “Sample” Cisco Nexus Switch Configurations

Nexus switch-A configuration

Upgrade the first Nexus switch with minimum System version 5.2(1)N1(3) (enter: show version)

Enable FCoE on the switch (disabled by default):

- configure terminal
- feature fcoe
- feature lacp

Note: The Nexus will require a reload.

Associate a VLAN with a VSAN for FCoE traffic:

- vlan 200
- fcoe vsan 200

Create a port channel with eth1/5 and eth1/6 with the same LACP Timer as defined by default in the Virtual Connect Domain, eth1/5 and eth1/6 can be used because both are mapped to the same UPC (i.e. UPC#1) on Nexus:

- interface Ethernet1/5

- description FCoE uplink to FlexFabric
- channel-group 200 mode active
- lacp rate fast

- interface Ethernet1/6
- description FCoE uplink to FlexFabric
- channel-group 200 mode active
- lacp rate fast

Create the VLANs for the IP traffic:

- vlan 1,10,20

- Create a trunk on the port channel interface to pass the FCoE (VLAN 200) and Ethernet traffic (VLAN 1,10,20):

  - interface port-channel 200
  - switchport mode trunk
  - switchport trunk allowed vlan 1,10,20,200
  - spanning-tree port type edge trunk

  Create a virtual Fibre Channel interface to carry the FCoE traffic through the port channel:

  - interface vfc 2005
  - bind interface port-channel 200
  - no shutdown

  Assign the vfc interface and the Fibre Channel port to the FCoE VSAN:

  - vsan database
  - vsan 200
  - vsan 200 interface fc2/1
  - vsan 200 interface vfc 2005

  Configure the interface connected to the datacenter LAN:

  - interface eth 1/17
  - switchport mode trunk
  - switchport trunk allowed vlan 1,10,20

  Configuration of the zone:

  Create zones:
  - zone name fcoe-zone vsan 200
  - member pwwn 21:53:00:02:ac:00:15:9d {This is the WWN of the first 3PAR controller port}
  - member pwwn 50:06:0b:00:00:c3:1a:20 {This is the WWN of the Blade FlexFabric Adapter port

Create zoneset:

- zoneset name zoneset1 vsan 200
- member fcoe-zone

Activate zoneset:

- zoneset activate name zoneset1 vsan 200

Nexus switch-B configuration

Upgrade the second Nexus switch with minimum System version 5.2(1)N1(3) (enter: show version)

Enable FCoE on the switch (disabled by default):

- configure terminal
- feature fcoe
- feature lacp

Note: The Nexus will require a reload.

Associate a VLAN with a VSAN for FCoE traffic:

- vlan 201
- fcoe vsan 201

Create a port channel with eth1/5 and eth1/6 with the same LACP Timer as defined by default in the Virtual Connect Domain, eth1/5 and eth1/6 can be used because both are mapped to the same UPC (i.e. UPC#1) on Nexus:

- interface Ethernet1/5
- channel-group 200 mode active
- lacp rate fast

- interface Ethernet1/6
- channel-group 200 mode active
- lacp rate fast

Create the VLANs for the IP traffic:

- vlan 1,10,20

Create a trunk on the port channel interface to pass the FCoE (VLAN 200) and Ethernet traffic (VLAN 1,10,20):

- interface port-channel 200
- switchport mode trunk
- switchport trunk allowed vlan 1,10,20,201
- spanning-tree port type edge trunk

Create a virtual Fibre Channel interface to carry the FCoE traffic through eth1/5:

- interface vfc 2005
- bind interface port-channel 200
- no shutdown

Assign the vfc interface and the Fibre Channel port to the FCoE VSAN:

- vsan database
- vsan 201
- vsan 201 interface fc 2/1
- vsan 201 interface vfc 2005

Configure the interface connected to the datacenter LAN:

- interface eth 1/17
- switchport mode trunk
- switchport trunk allowed vlan 1,10,20

Configuration of the zone:

Create zones:

- zone name fcoe-zone vsan 201
- member pwwn 20:53:00:02:ac:00:15:9d {This is the WWN of the second 3PAR controller port}
- member pwwn 50:06:0b:00:00:c3:1a:22 {This is the WWN of the Blade FlexFabric Adapter port 2}

Create zoneset:

- zoneset name zoneset1 vsan 201
- member fcoe-zone

Activate zoneset:

- zoneset activate name zoneset1 vsan 201

Comparing VC and VMware Networking Technology

One method of understanding how Virtual Connect operates on the LAN is to compare the Virtual Connect networking components and their functionality to the networking components of a VMware ESX server. Since the networking technology used in a VMware ESX server is commonly understood and accepted by many customers, understanding the many similarities between VC and a VMware ESX server will help an implementer have a better understanding of how Virtual Connect looks to, and behaves on, the external network. Just to be clear, Virtual Connect and a VMware ESX server are fundamentally different products and address completely different needs within the datacenter. This comparison is strictly about understanding the similarities between the two products

in regards to networking technology in order to better understand Virtual Connect.

A Description of the VMware Components:

Referencing the figure below, the VMware ESX server (left) is a single physical server running server virtualization software (VMware) that allows the physical server to host one or more instances of a virtual server, called a Virtual Machine (VM). In addition, the VMware ESX server provides external network connectivity to the internal servers (VMs) using a virtual (software) implementation of a layer 2 bridge, called a vSwitch. The VM virtual NICs (vNics) are assigned to one of the vSwitches and the vSwitches are then associated with real physical NICs residing in I/O slots on the VMware ESX server. The vSwitches can have one or more physical NICs (uplinks) assigned to them to provide external network connectivity. If more than one physical NIC is assigned to the same vSwitch, network redundancy and/or load balancing is provided for the internal servers (VMs) assigned to that vSwitch. The physical NICs then present one or more MAC addresses to the external network, depending on the number of VMs communicating to the external network through each physical
NIC.

A Comparative Description of the VC Components:

Referencing the figure below, the c-Class enclosure (right) is a single physical enclosure that hosts one or more real physical servers, called a blade server. In addition, the c-Class enclosure provides external network connectivity to the internal servers (blade servers) using a hardware implementation of a layer 2 bridge, called a Virtual Connect Ethernet network (vNet). The blade server physical NICs (pNics) are assigned to one of the vNets and the vNets are then associated with real physical VC uplink ports from VC-Enet modules residing in the I/O bays on the c-Class enclosure. The vNets can have one or more VC uplinks assigned to them to provide external network connectivity. If more than one VC uplink is assigned to the same vNet, network redundancy and/or load balancing is provided for the internal servers (blade servers) assigned to that vNet. The VC uplinks then present one or more MAC addresses to the external network, depending on the number of blade servers communicating to the external network through each VC uplink.

Figure 18. VMware ESX server Compared to Virtual Connect Enclosure

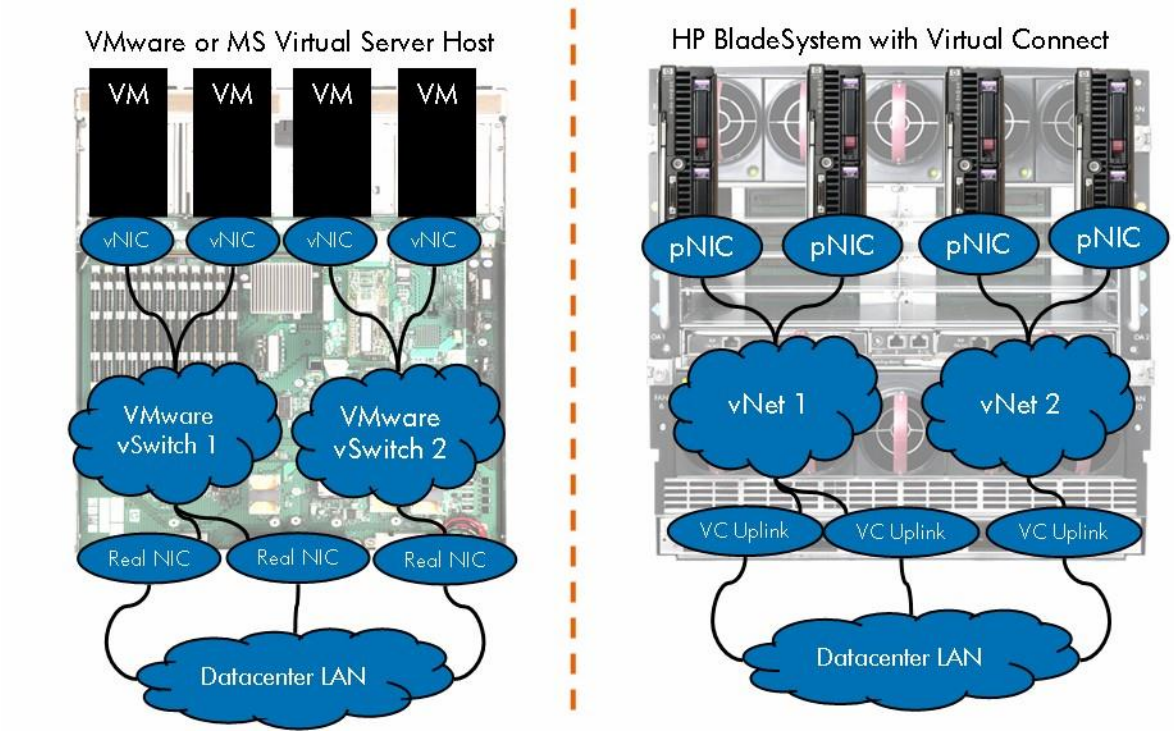

After comparing the components and their functionality, it is obvious why many customers treat a c-Class enclosure with Virtual Connect the same way they would a single VMware ESX server. In other words, VC allows an entire enclosure to look to the network just like a big VMware ESX server. From a network redundancy and load balancing perspective, from a security perspective, from a port monitoring perspective, etc.; VC simplifies the network connectivity for an entire enclosure and makes it behave on the network like a single VMware ESX server.

Just to reiterate the commonality between the two products from a networking technology perspective, below is a table of similarities.

Figure 19: Networking Technology Similarities Between VC and VMware

	VMware	Virtual Connect
Provides network connectivity for multiple servers	(virtual)	(physical & virtual)
Provides redundancy external connectivity for internal servers		
Provides load balanced external connectivity for internal servers		

External network sees multiple MAC addresses on pNICs (VMware) or VC	(from VMs)	(from blade servers)
Can be configured to isolate internal servers into separate Layer 2 domains (broadcast domains)		
Can be configured to allow internal servers to communicate directly		
pNICs\VC Uplinks can be configured as Port Trunks (EtherChannel)		
pNICs\VC Uplinks can be configured as VLAN Trunks		
Does NOT transmits BPDUs to external network		
Does NOT participates in the Datacenter Spanning Tree	(uses internal loop prevention)	(uses internal loop prevention)
Uses a virtual (software) component to pass frames between internal servers and external network		X
Uses a hardware component to pass frames between internal servers and external network	X	
Is Operating System Independent	X	
Supports both Itanium and x86 platforms	X	

Virtual Connect's Loop Prevention Technology

Virtual Connect's loop prevention technology is best described by building on what was discussed in the section above, "Comparing VC and VMware Networking Technology". VC's loop prevention technology is very similar to the NIC Teaming\bonding technology used by ProLiant servers.

Referring back to the figure above comparing VMware and VC, it can be seen that the ESX server is using NIC bonding technology to provide load balancing and\or fault tolerance for vSwitch 1 and

for the associated VMs. Even though vSwitch 1 is redundantly connected to the Data Center LAN, no loops are formed. In addition, Spanning Tree is not needed between the ESX server and the Data Center LAN to prevent the loop. Instead, the NIC bonding technology on the ESX host prevents loops on the network by only allowing one logical path (single NIC port or single channel group\port trunk) to be active at one time. In comparison, VC uplinks prevent loops in the same manner. In other words, on a per Virtual Connect network (vNet) basis, each vNet prevents loops on the Data Center LAN by only allowing one active logical path at a time.

The behavior of the VC uplink ports associated with a vNet can be compared to the NIC ports associated with a vSwitch. In both cases, Spanning Tree is not needed to prevent loops on the Data Center LAN. Instead, in both cases "port bonding" technology is used to only allow a single active logical path at any given time.

Stacking Virtual Connect Ethernet Modules

Virtual Connect can be configured by an Administrator to allow any VC uplink on any VC Ethernet module to provide external connectivity for any server downlink. VC can also be configured to allow connectivity between any set of server downlinks on any VC Ethernet module. Virtual Connect provides this flexible connectivity through the use of 'stacking links' between VC Ethernet modules.

Stacking links allow all VC Ethernet modules in the VC Domain to be configured as, and operate as, a single device (except for port channeling). Stacking links are to a group of VC Ethernet modules what the PCI bus is for a team of server NICs – a common communication path allowing all devices to work together as one.

When any two Virtual Connect Ethernet modules from the same VC domain are directly connected together using 1 Gb or 10 Gb ports, the two modules automatically negotiate, using LLDP, the link as a stacking link. No manual configuration is required to make an uplink a stacking link. Simply connecting two VC uplinks together is all that is required.

In the figure below, the green horizontal lines represent the internal stacking links that exist between horizontally-adjacent VC modules by default. The orange vertical lines represent the external stacking links that HP recommends customer install using 10 Gb or 1 Gb ports. The minimum recommendations below provide stacking link redundancy for each VC Ethernet module.

Note:
Stacking links are required for all Virtual Connect Ethernet modules in the same Virtual Connect Domain.

Figure 20. VC Stacking Link Recommendations

Two Enclosure

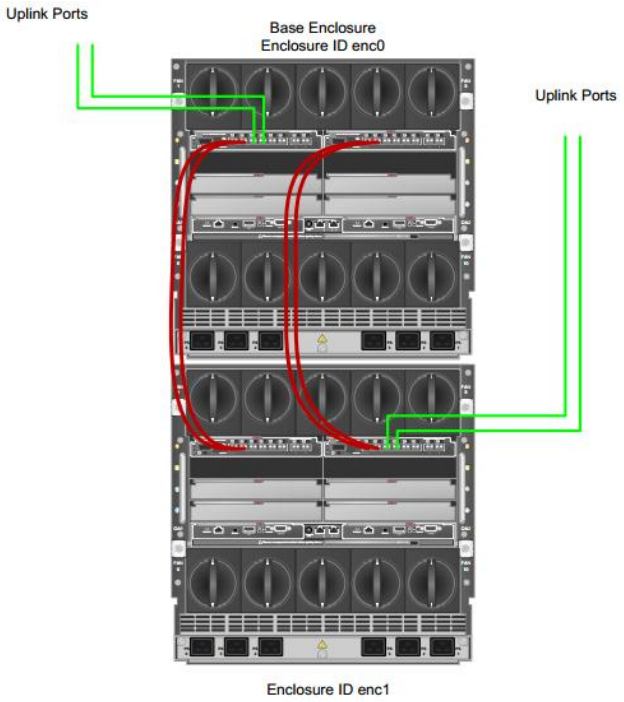

Three Enclosures

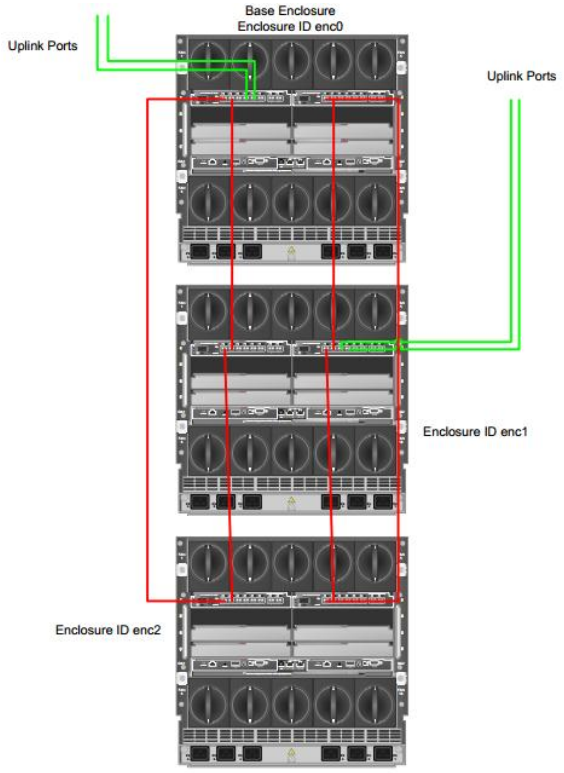

Four Enclosures:

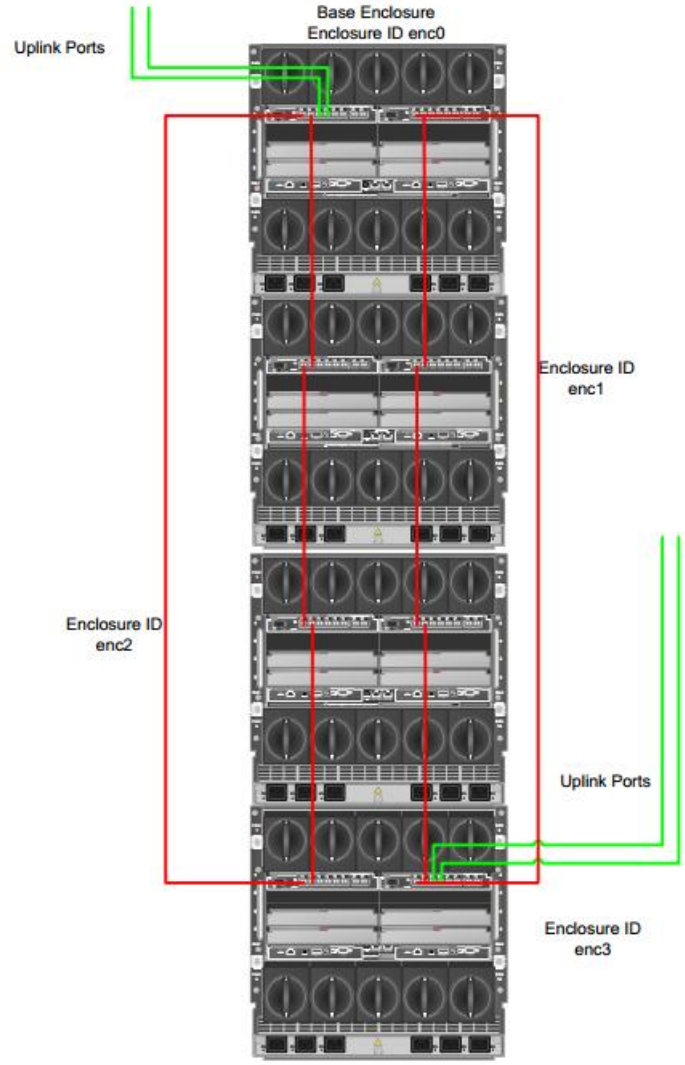

Not Recommended:

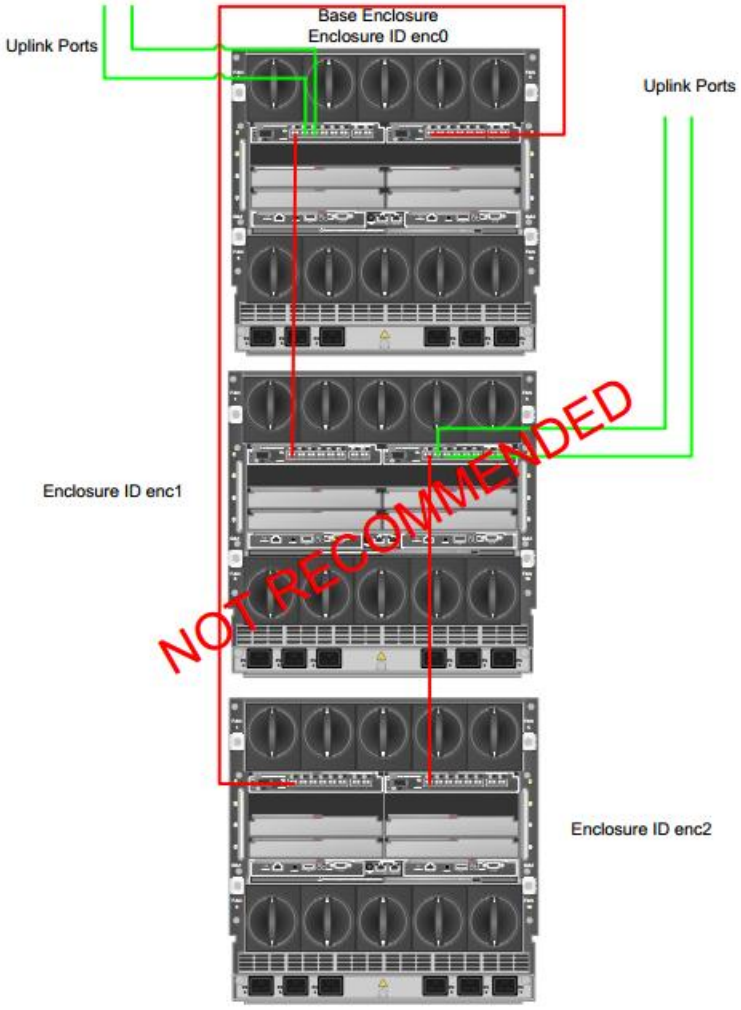

Optimizing Virtual Connect Stacking Links

Virtual Connect stacking links provide the physical path between downlinks (server NIC ports) to VC uplinks. As a result, the fewer VC uplinks a frame has to traverse, the less latency the frame incurs in reaching the external network. Each VC module a frame must traverse adds approximately 3.8 microseconds of latency to the frame. Because of the additional latency that each hop can add, an administrator may wish to add additional stacking links above the recommended minimum.

On a per-vNet or per-Shared Uplink Set basis, the VC-Enet module with the active uplink, or port channel, is the module that all server frames will use to exit the VC domain towards the external network. For example, if a server NIC port (assigned to vNet1) is connected to VC-Enet module 2 and the active uplink for vNet1 is connected to VC-Enet module 6, then module 2 will send the frame towards module 8 via the shortest stacking link path. If the modules are stacked together using stacking links as depicted in the figure above, module 2 will take the path to module eight via the stacking link to module 7. Module 7 will then use it's directly connected internal stacking link to deliver the frame to module 8 and out the uplink towards the external network. Module 2 could use the longer path via module 1 to module 3 to module 4 to module 8 if it loses its shorter path through module 7.

An administrator could purposefully add additional stacking links to provide shorter stacking link paths between modules. For example using the scenario above, the administrator could add an additional stacking link between module 2 and module 8. If added, module 2 could send frames towards the active uplink on module 8 using the directly connected stacking link to module 8. Purposefully adding additional stacking links is not a common practice since each module's latency of 3.8 microseconds is so low.

VC Managed MAC Addresses

One of the many features provided by Virtual Connect is the ability to "manage" the blade server MAC addresses. Specifically, Virtual Connect 'manages' the blade server MAC addresses. Virtual Connect does not 'virtualize' the blade server MAC addresses. Many VC implementers don't appreciate the difference between 'virtualized' MAC addresses and 'managed' MAC addresses.

A 'virtualized' MAC address is a MAC address that is not really owned and used by a physical NIC. Often, a virtualized MAC address is a MAC address that replaces the real MAC address of a physical NIC without the server's knowledge. In other words, the server thinks it is communicating on the network with MAC address X, however, some device is replacing real MAC address X with a virtual MAC address Y. Many of the benefits of MAC address management are lost in this type of implementation.

A 'managed' MAC address is a MAC address that really is owned and used by a physical NIC. The server has simply been assigned, by the administrator, to use a specific MAC address on a specific physical NIC port. This managed MAC address, for all intents and purposes, appears to the server as the MAC address that was burned into the physical NIC at the factory.

The benefits of VC Managed MAC addresses are:

- **WYSIWYG - What You See (on the server) Is What You Get (on the network)**
  There is no discrepancy between what the server thinks its MAC address is and what the external network sees as the server's MAC address. This dramatically reduces the

complexity of troubleshooting network related issues.

- **Server application licensing is maintained after hardware changes**
  Many server application licensing mechanisms can key off the server's MAC addresses. If the server's MAC address changes (replacing a failed NIC, booting server image on a different physical server, etc.), then the application licensing may require re-licensing using the new MAC address. Virtualized MAC addresses do not prevent this problem. However, VC's use of managed MAC addresses does prevent this problem since the server image will always see the VC managed MAC address regardless of which physical server it is running
  on.

- **No Performance impact on network devices**
  Virtualized MAC addresses can require that a network device (e.g. switch) manipulate every frame a server transmits to replace the server's MAC address with the virtualized
  MAC address. Also, when the source MAC address is edited by the network device, the frames checksum (CRC) has to be recomputed by the network device. The more frames a server transmits, the more work the network device has to do, which can have an impact on the performance of the network device. Alternatively, VC's use of managed MAC addresses means the server transmits with the managed MAC address. No device on the network (VC or switch) is required to manipulate the server's frames. This results in absolutely no performance impact on the network.

When a blade server with VC managed MAC addresses is removed from the enclosure (or when a VC Server Profile is unassigned from a server blade while still in the enclosure) the blade server automatically reverts back to the actual MAC addresses burned into the physical NICs at the factory. This prevents any issues with duplicate MAC addresses on the network caused by moving blade servers around within the data center.

Virtual Connect Manager (VCM) provides three domain-wide choices for managing the blade server MAC addresses:

- **Static, factory-default MAC addresses**
  As the name suggests, this setting tells Virtual Connect to not manage the server MAC addresses. The server will only use the original factory burned-in MAC address.

- **User-defined**
  This setting allows the Administrator to define a Locally Administered MAC address range that Virtual Connect will use to assign to blade servers.

- **HP Pre-defined** (recommended)
  This setting tells Virtual Connect to assign server MAC addresses from the pool of MAC addresses that HP has reserved. This option allows the Administrator to choose from one
  of 64 ranges of 1024 MAC addresses.

In addition, when using either the User-defined or HP Pre-defined settings as the domain-wide default setting, each individual VC Server Profile can optionally override the domain-wide setting by selecting to use the factory-default MAC address instead.

Important points about VC managed MAC addresses:

- VC only manages the MAC address of physical NIC ports on a blade server. VC does not manage the MAC address used by virtual servers that may run on a physical server.

- A server reverts to the factory-default MAC address when removed from the enclosure or when no VC Server Profile is applied to it.
- VC generates 2 MAC addresses per NIC port in order to support multi-function adapters that may require more than one MAC address per port.
- Customers requiring a centralized database for managing the VC managed MAC address pools should consider implementing Virtual Connect Enterprise Manager. See www.hp.com/go/vcem for more details.

Spanning Tree

Since Virtual Connect does not present itself to the external Cisco network as another "Ethernet switch", the subject of "spanning tree interoperability" is not an applicable topic when discussing the two products. In other words, in the same way that a server (with multiple NICs) does not need to be configured for spanning tree, Virtual Connect (with multiple uplinks that look like NICs to the external Cisco network) does not need to be configured for spanning tree either. This means that the network administrator does not have to worry about STP interoperability issues between VC and Cisco since VC uplink ports do not transmit or receive STP BPDUs (Spanning Tree configuration frames) and cannot interfere with the external switched infrastructure's implementation of STP any more than a multihomed server can. This also means that all questions about PVST+ support, Rapid Spanning Tree, MST, etc. are not applicable for Virtual Connect deployment discussions.

Since VC uplink ports simply look to the network like "regular old server ports" with multiple MAC addresses, HP recommends that Cisco downlink ports connected to the VC uplink ports be configured as if they are connected directly to server NIC ports. In other words, configure the Cisco downlink ports as "access" or "trunk" ports with PortFast enabled. Enabling BPDUGuard is also an acceptable option for customers that regularly enable it for other Cisco switch ports connected directly to server ports. In summary, since Virtual Connect allows the c-Class enclosure to connect to the external network in much the same way as a VMware ESX server connects to the external network, it is easy to understand why Virtual Connect, just like a VMware ESX server, does not participate in the external network's spanning tree domain.

Unidirectional Link Detection (UDLD)

UDLD is a Cisco proprietary feature that only operates between two directly connected Cisco devices. The purpose of UDLD is to detect unidirectional links (UniDirectional Link Detection) in situations where spanning tree is being used to detect and prevent loops. If a unidirectional link goes undetected, spanning tree won't properly detect the loop and won't block the port to prevent a loop on the network. Since spanning tree is not used between VC uplinks and external Cisco switches, a spanning tree loop caused by a unidirectional link is not possible. Therefore, UDLD is not required nor supported on Virtual Connect uplinks.

Quality of Service

In the current version of firmware, Virtual Connect does support user-configurable settings for traffic classification, marking, or prioritization. VC does utilize layer 2 quality of service markings (802.1p Class of Service) and layer 3 quality of service markings (TOS, DSCP) for frame prioritization. There are three modes of operation, 1. No QOS – First in First out. 2. Honoring mode, VC will honor the markings already made

If an Administrator needs the blade server NIC ports to receive layer 2 QoS markings (Class of Service), then the server NIC ports will need to be assigned to a vNet operating in 802.1Q\p tunneling mode. See the section above entitled "VC Uplinks and VLAN Trunking" for more information.

Security

ACLs & VLAN ACLs

In the current version of firmware, Virtual Connect does not support user-configurable port or VLAN ACLs. However, an Administrator can specifically design the deployment of vNets to make use of ACLs or VACLs configured on the upstream Cisco switch.

In the figure below as an example, all four blade servers are on VLAN 1 even though they are split between two different vNets. Since the uplink ports for both vNets are connected to Cisco switch ports assigned to VLAN 1, both vNets represent VLAN 1 within the VC Domain. While server 1 can communicate directly with Server 2 without leaving the VC Domain, Server 1 cannot talk to Servers 3 or 4 without transiting the external Cisco switch. Server 1's traffic will exit the VC domain via vNet VC_LAN1_A connected to interface gi0/1, be switched by the Cisco switch over to interface gi0/3, and reach Server 3 via the vNet VC_LAN1_B.

Since Server 1's traffic must transit the external Cisco switch to reach either Server 3 or Server 4, Server 1's traffic will be subject to any port ACLs or VLAN ACLs configured on the external Cisco switch. In this case, the Administrator could apply a VACL on VLAN 1 of the external Cisco switch to prevent Server 1 from communicating with Server 3 while still allowing Server 1 to communicate with Server 4.

The scenario above is just one example of the many ways an Administrator can creatively use vNets to accommodate complex network configurations.

Figure 21: Using multiple vNets to force server-to-server traffic through external Cisco switch

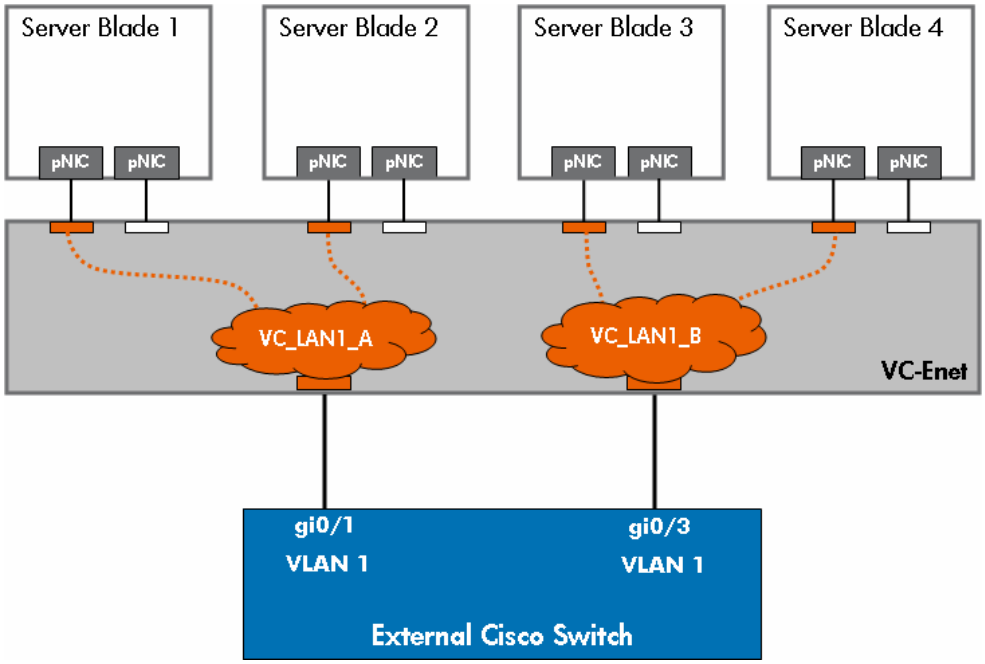

Port Security

Many network administrators use a Cisco switch feature called "Port Security" to provide additional security on the network. This feature allows the administrator to control how many MAC address are learned on a particular switch port or allows the administrator to limit connectivity to specific MAC addresses. Typically, the feature is enabled on the Cisco switch ports where desktops and laptops are connected in order to prevent additional hubs or switches from being connected to the network. By enabling this feature, the administrator can configure the port to automatically shut off (err-disable) if more than a certain number of MAC addresses are learned on the switch port. When an unauthorized switch or hub is connected to the port, it is possible that more MAC addresses are learned than are permitted by the administrator and the switch port is shut down. This isolates the unauthorized switch or hub until the administrator re-enables the port.

Port security can also be used in the data center for Cisco switch ports connected to server NICs. Since a server NIC port typically only uses a single MAC address, the feature does not cause the port to be shut down. However, if Port Security is enabled on the Cisco switch ports connected to VC uplinks and the administrator limits the number of MAC addresses to less than the number of MAC addresses being used by the servers in the BladeSystem enclosure, then the port security feature can cause the Cisco switch ports to shut down. This causes the BladeSystem enclosure to be isolated from the external network.

If port security must be used on the Cisco switch ports connected to Virtual Connect, make sure the number of MAC addresses allowed is greater than or equal to the number of MAC addresses used within the enclosure. Be sure to take into account servers that may be added to the enclosure at a later date and take into account the virtual MAC addresses used by virtual servers (e.g. ESX VMs).

If port security is configured to only allow specific MAC addresses to communicate on the Cisco switch port connected to the VC Uplink, an Administrator may find that using HP Managed MAC

addresses with a user-defined range will simply the task. Simply configure port security to allow the same range of MAC addresses that are manually configured for the user-defined range.

Whether an administrator is configuring port security to allow a certain number of MAC address or to allow specific MAC addresses, they must configure all Cisco ports assigned to the same vNet (or Shared Uplink Set) with the same port security settings in order to eliminate communication problems after a VC uplink failover.

Private VLANs

Private VLANs, or PVLANs, is a Cisco feature that allows switch ports assigned to the same VLAN to be configured for layer 2 isolation from other switch ports in the same VLAN. The benefit of this type of feature is enhanced security. For example, an Administrator can assign several servers to the same VLAN, utilizing the same IP subnet, and only allow the servers to directly communicate with the default gateway but not with each other.

In the current release of VC firmware, Virtual Connect does not implement support for Private VLANs within a single vNet. However, an Administrator can achieve complete Layer 2 separation between server ports by assigning the server ports to separate vNets. Each vNet is a self-contained Layer 2 network\broadcast domain. Virtual Connect does not internally permit frames to pass between vNets. This means that a customer can use multiple vNets to isolate servers from each other within VC.

An Administrator can use multiple vNets to extend the function of PVLANs configured on the upstream Cisco switch. In the example below, two VC vNets and a Cisco Private VLAN have been used to create two “community” groups and one promiscuous port. All four servers are on the same VLAN (VLAN 1), however, they are segregated at Layer 2 within VC because they are assigned to two different vNets (VC_LAN1_A & VC_LAN1_B). These two vNets are uplinked into a Cisco switch where VLAN 1 has been configured as a Private VLAN. Interface gi0/1 connected to VC_LAN1_A is an “isolated” port, so it only permits VC_LAN1_A to communicate with the router port (promiscuous). Also, VC_LAN1_B is connected to interface gi0/3 configured as “isolated” for VLAN 1, so it can also only talk to the router. Server 1 and Server 2 can directly communicate with each other within the VC Domain, however, they both would have to transit the external Cisco switch in order to communicate with Server 3 and Server 4. As a result, all communication from Server 1 or 2 to Server 3 or 4 is subject to the PVLAN configuration on the external switch. If gi0/1 and gi0/3 are isolated ports and gi0/5 is a promiscuous port, then Server 1 and Server 2 can communicate with each other and with the router, but they cannot communicate with Server 3 or 4. The same is true for Server 3 and 4 – they can communicate with each other and the router, but not Servers 1 and 2. Both VC_LAN1_A and VC_LAN1_B function as Community VLANs with the PVLAN. Note: all servers are on the same VLAN (VLAN 1) and can be assigned IP addresses from the same subnet.

Figure 22: Using multiple vNets to extend PVLAN configuration from external Cisco switch

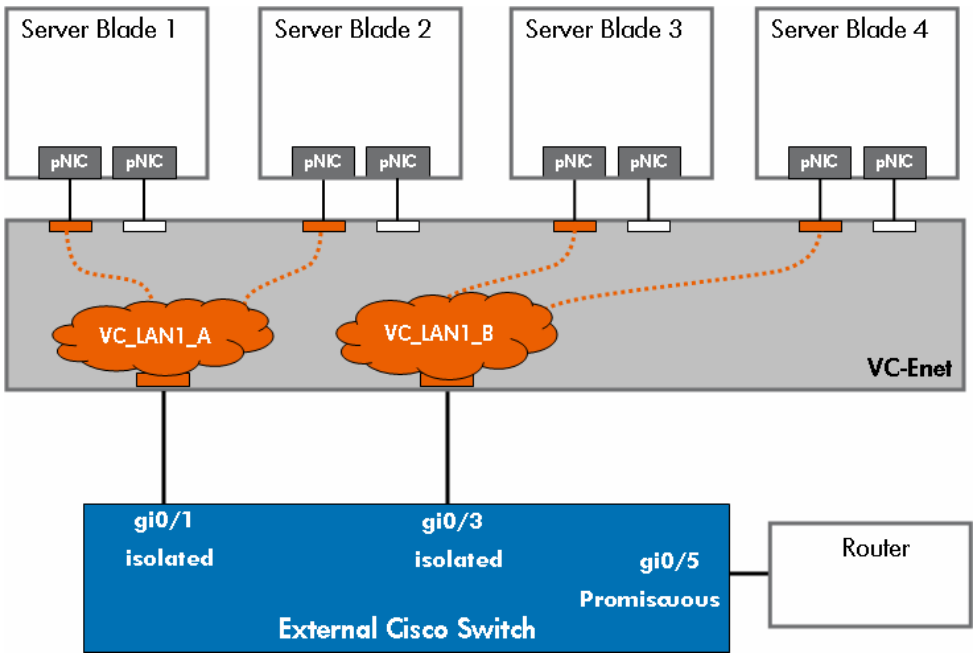

(Please refer to Cisco documentation for a discussion of isolated, promiscuous, and community VLANs)

Multicast & IGMP Snooping

The IGMP Snooping feature allows VC-Enet modules to monitor (snoop) the IGMP membership activities of the blade servers and optimize a vNet's handling of multicast traffic to maximize network resource utilization. Currently only IGMP v1 and v2 (RFC2236) are supported.

The IGMP Snooping idle timeout interval is set to 260 seconds by default. This value is basically the "Group Membership Interval" value as specified by IGMP v2 specification (RFC2236). For optimum network resource utilization, set the interval to match the configuration on the network's multicast router.

By default, VC will flood, not prune, for all multicast groups. If a particular multicast group receives an explicit join from a server NIC port, then that multicast group (and that group only) will be pruned for any server NIC port that doesn't transmit a join message. All other multicast groups will continue to flood.

See the Virtual Connect User Guide for instructions on configuring IGMP Snooping.

SPAN (Port Monitoring)

Virtual Connect supports port monitoring functionality to assist in troubleshooting networking issues for servers connected to the external network through Virtual Connect. VC allows an Administrator to define a single, active port monitoring session per VC domain.

The monitor session must be configured with at least one 'monitored port' and a single 'analyzer port'. The 'monitor port' list is the list of server downlinks whose traffic will be mirrored. The 'analyzer port' is the VC uplink port that the network analyzer is connected to. VC will mirror the traffic from the monitored ports to the analyzer port.

A Port Monitoring session can mirror the traffic for up to 16 server downlinks to the analyzer port. Any VC uplink can be defined as the analyzer port and any server downlinks can be selected as the monitored port. There is no VC Ethernet module dependency.

Lastly, the Administrator can choose the direction of the traffic to mirror – to the server, from the server, or both.

Logging

Virtual Connect provides a local system log that records VC domain events and timestamps. The local system log is not editable and does not permit manual clearing (for security purposes). The log is circular and once the buffer is full, older entries are removed. Virtual connect logs can also be redirected out to an external syslog server.

Centralized User Access Control (LDAP)

In the current version of VC firmware, TACACS+ and RADIUS are also supported for centralized user access control.

Please consult the Virtual Connect User Guide for LDAP, TACACS+ and RADIUS implementation details.

Cisco Discovery Protocol (CDP)

Cisco Discovery Protocol (CDP) is a Cisco proprietary protocol used to exchange device identity and capability information between two directly connected Cisco devices. CDP allows each Cisco device to display, for the administrator, a table showing the identity and capability of other directly connected Cisco devices. One of the primary uses of this information is for troubleshooting purposes.

Virtual Connect does not support CDP. VC does support the industry standard protocol called Link Layer Discovery Protocol (LLDP) by default. LLDP is functionally equivalent to CDP, although the two protocols are not compatible. Many Cisco switches support LLDP. The global configuration command "lldp run" and the interface commands "lldp transmit" and "lldp receive" are used to enable LLDP on a Cisco switch that supports the protocol. The exec command "show lldp neighbors" displays LLDP neighbor information to the Cisco switch console.

Virtual Connect provides the LLDP neighbor information for each VC uplink at the bottom of the statistics page for each port.

Additional Resources

HP Services:
www.hp.com/go/bladesystem/services

BladeSystem Solutions:
www.hp.com/go/bladesystem/solutions

Virtual Connect Cookbook:
www.hp.com/go/bladeconnect (see the Virtual Connect Interest Group)

Virtual Connect Documentation:
www.hp.com/go/bladesystem/documentation

Virtual Connect Firmware:
www.hp.com/go/bladesystemupdates

HP NIC Teaming for Windows Whitepaper:
ftp://ftp.compaq.com/pub/products/servers/networking/TeamingWP.pdf

Appendixes

Appendix A: Description of VC Network Diagram

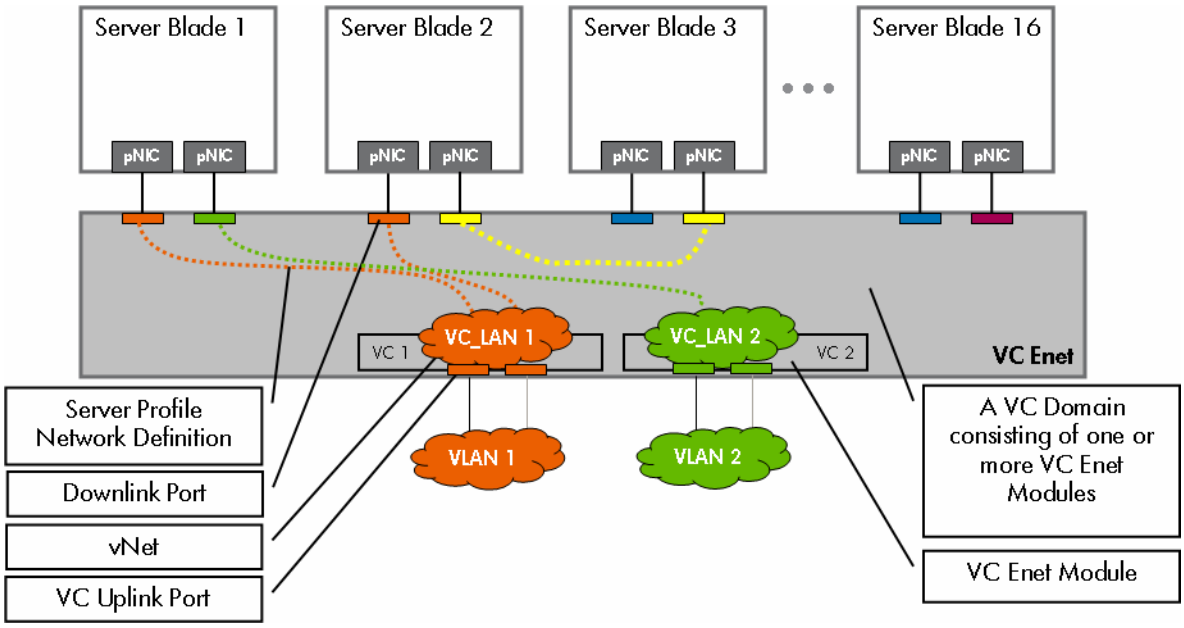

Appendix B: c3000 Port Mapping Diagram

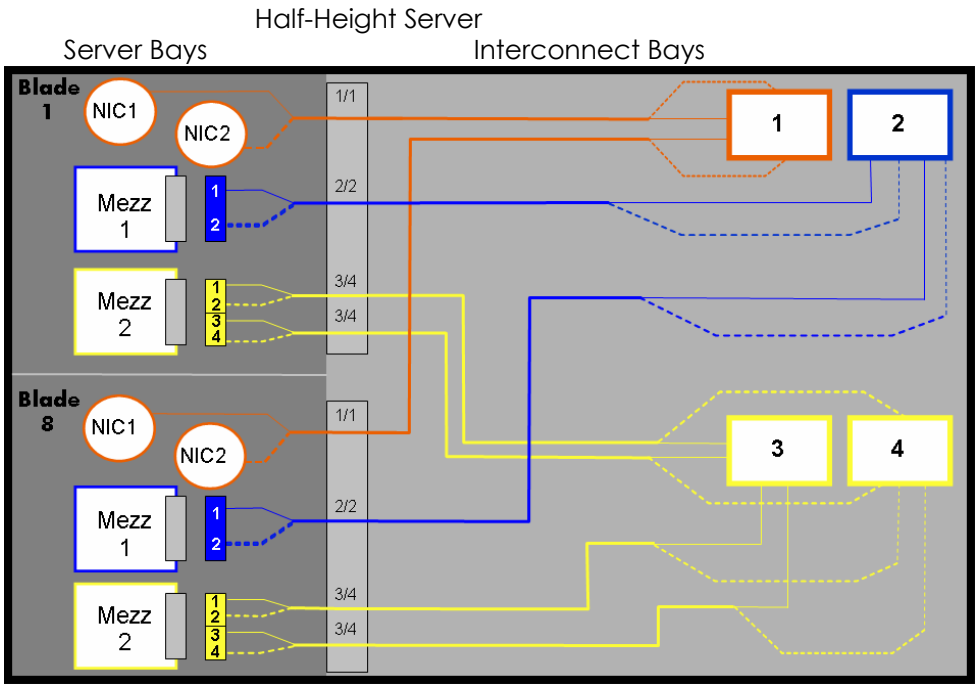

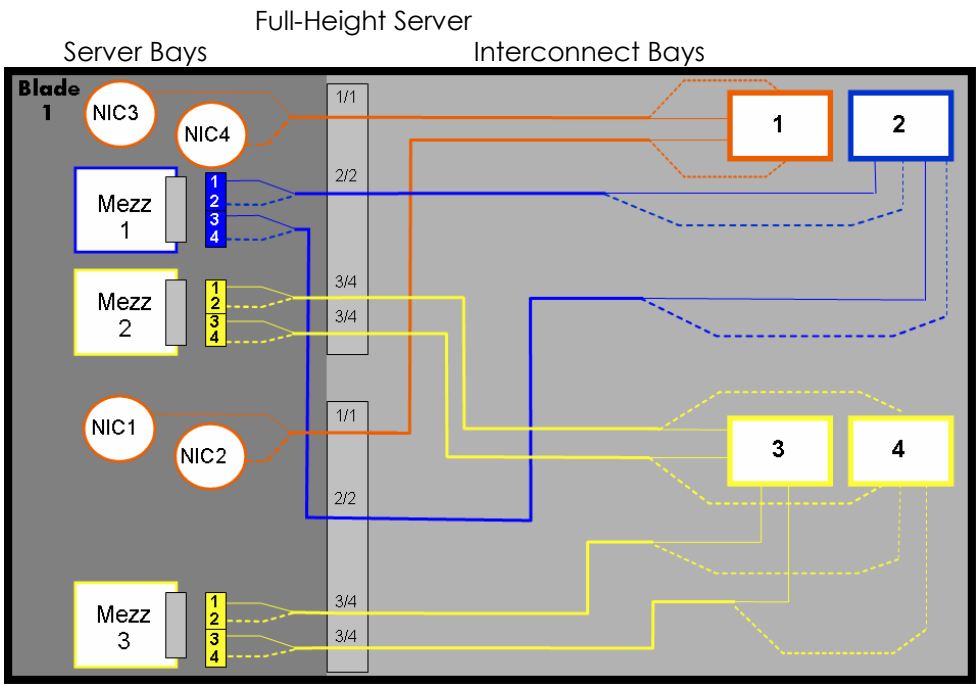

Appendix C: Frequently Asked Questions

Q1: Why do I see lots of dropped frames (discards) on standby VC uplink ports?
A1: An external switch has no concept of which VC link is the active uplink and which is the standby uplink. As far as the external switch is concerned, one of the uplinks is just a whole lot busier. That means that the external switch is still going to send some types of frames down the standby link and the standby link is going to discard them. This includes all broadcasts, multicasts, and unknown unicasts (destination lookup failure in the CAM table on the external switch). Any of these frames that are received on the standby link will be dropped and will be reflected in the counters.

Q2: Can I manually choose which port channel is the preferred channel for a vNet?
A2: No, VC does not currently support setting the 'port role' for vNets (or Shared Uplink Sets) with LACP enabled (connect mode 'auto'). VC determines which port channel is active and which is standby, based on the following criteria, 1 - Number of functional uplinks per port channel, 2 – Total bandwidth provided by each port channel, 3 – if the previous two are a tie, then the VC module with the lowest MAC address (see TOE tag) will provide the active port channel. With the information above, the Administrator could construct the primary port channel with one addition uplink in order to make it the preferred port channel.

Q3: Do I have to use the same load balancing algorithm on both sides of the same port channel?

A3: No, you can have different load balancing algorithms on each side of a port channel. But as a course of good network design it is best to use the same load balancing in each device in the path to insure the most optimal traffic flow.

Q4: I see Link Layer Discovery Protocol (LLDP) frames when I connect a network trace analyzer to a VC uplink port. What is VC doing with LLDP and can I disable it?
A4: LLDP is the IEEE equivalent to Cisco Discovery Protocol (CDP). It is a Layer 2 protocol that allows one device to both announce itself (and some of its feature set) to a neighboring device as well as discover other connected devices on the network. It is extremely low-bandwidth and is unobtrusive. VC uses LLDP to determine when one of its uplinks or cross-connects is directly connected to another VC module in the same VC domain so that it can form a stacking link. There is currently no way to disable it.

Q5: VC Ethernet is just another switch
A5: Incorrect: While VC uses tried and true, IEEE standard, Layer 2 bridging functionality, its primary purpose is to provide many server virtualization and management features that are nonexistent in traditional switches. VC may perform some functions like a traditional switch; however, VC has many additional features which clearly distinguish it from a traditional switch. Likewise, server virtualization hypervisors (for example, VMware ESX, Microsoft HyperV, Citrix Xen) perform some functions of a traditional switch but, like VC, have many additional features which clearly distinguish them from a traditional switch. As a result, it is incorrect to say that either technology, VC or hypervisor virtual switching, is "just another switch". VC and server virtualization hypervisors are very similar in the networking functionality that they provide to servers; a hypervisor provides it for virtual servers and VC provides it for physical HP server blades. In the same way that a hypervisor provides this functionality in a way that interoperates with the external network, VC also provides this interoperable connectivity between HP server blades and the external network. Virtual Connect is not called a "switch" because it is not configured, deployed, or managed as a switch and does not present itself to the external network as a switch again, much like a hypervisor. When Virtual Connect is linked to the external network, the external network "sees" the same behavior from VC as it "sees" when a server hosting a hypervisor is connected to the external network. Since VC is not configured, deployed, or managed like a traditional switch and presents itself to the network as an endpoint (like a server), it is incorrect to call VC a "switch".

Q6: I am trying to get 802.3ad Port Trunking to work but can't seem to get it to pass traffic. What am I doing wrong?
A6: VC currently only supports LACP for 802.3ad port trunks. Cisco's PAgP is not supported by VC. Future versions of VC may add additional features to enhance this functionality. For LACP to work properly on a Cisco switch, the channel mode must be set to either Active or Passive. VC can display detailed Ethernet and port trunk statistics by clicking the desired Ethernet module under Hardware Overview in the left hand tree-view of VCM.

Q7: How do I setup a cluster heartbeat network in Virtual Connect?
A7: Create a vNet in VC and do not assign a VC uplink port to it. Next, assign a "heartbeat" NIC from each blade in the cluster to this vNet. All heartbeat traffic will be contained within the vNet and will not be transmitted outside of the enclosure.

Q8: I need more than 16 VC uplinks. If I add more VC Ethernet modules to add more uplinks, am I required to use additional NICs on my servers?
A8: No, you can add more VC Ethernet modules, stack them with the other VC Ethernet modules and just use the uplink ports. Any VC uplink on any VC Ethernet module can be used to provide external connectivity for any downlink on any VC Ethernet module.

Q9: I need more NICs on my blade servers. If I add more VC Ethernet modules to add more downlink ports, am I required to use additional VC uplinks ports to provide connectivity for these new downlink ports?
A9: No, you can add more VC Ethernet modules, stack them with the other VC Ethernet modules and the new downlink ports can be configured to use the uplinks on the existing VC-Enet modules. Any VC uplink on any VC Ethernet module can be used to provide external connectivity for any downlink on any VC Ethernet module.

Q10: I noticed that the VC Ethernet module in interconnect bay 1 is the active Virtual Connect Manager and that the VC module in bay 2 is the standby. Does this mean that only the VC module in bay 1 is providing Ethernet connectivity for the blade servers?

A10: No. Regardless of which VC module is running the active Virtual Connect Manager, all VC modules can be used simultaneously to provide network connectivity.

Q11: Does VC support iSCSI?
A11: Yes VC is compatible with iSCSI. Since VC is a layer 2 device and iSCSI is an upper layer protocol, above TCP/IP, VC does not implement any features specific to iSCSI. However, VC can provide network connectivity for a server running iSCSI just like any other protocol.

Q12: Why are failovers taking longer than expected to restore connectivity for the blade servers (taking 20 seconds or longer)?
A12: Make sure the upstream Cisco switch ports connected to the VC uplink ports are configured with PortFast enabled ("spanning-tree portfast" or "spanning-tree portfast trunk")

Q13: Does VC interact with STP on my network?
A13: No. VC uplink ports look just like server NIC ports (ex. Physical NIC ports on an ESX server) and VC does not support STP on the VC uplink ports.

Q14: Should I expect BPDUs to be sent from the VC uplink ports to my external Cisco switch ports?

A14: No. VC uplink ports do not transmit BPDUs.

Q15: Can I extend any of my L3 routing protocols through the VC domain?
A15: Since VC is a layer 2 device, it does not support any routing protocols. However, layer 3 routing protocols such as OSPF, RIP, RIP2, etc. can be used on the servers and operate transparently through VC.

Q16: Can I configure transmit and receive load balancing NIC teaming with full redundancy throughout the VC domain?
A16: Yes, only if you are using our Integrity blades with INP for Windows/Linux, APA for HP-UX, or Smart Load Balancing in Linux on x86. There are currently no solutions for Windows on x86.

Q17: Does Virtual Connect Support QOS
A17: VC firmware 4.01 supports priority QoS queues and is available on all VC modules compatible with this version.
The implementation of QoS and CoS in the VC module complements the enforcement of prioritization at Flex-10 based adapters. VCM can allocate bandwidth based on up to 8traffic classes. VC will queue frames aligned to these classes, and then a VC scheduler will prioritize the traffic on both egress and ingress. Up to 8classes of service can be activated from the following selections:
•2 Fixed –best effort and lossless(FCoE)
•2 Predefined –real time and medium (latter not enabled at this time)
•4 classes that can be user-defined

Q18: Does VC support VMware and other OSs that support host-based VLAN tagging?

A18: Yes. See appropriate sections above.

Q19: Can I use third party branded SFPs and XFPs in a VC 1/10 Gb-F module?

A19: No, only HP branded SFP and XFP modules are supported.

Q20: Will upgrading the VC firmware require an outage?
A20: In general, HP recommends upgrading VC firmware during a scheduled maintenance window. However, VC is able to perform a rolling, non-intrusive upgrade of all modules so long as redundancy is configured throughout the solution.

Q21: How fast should I expect a failed VC uplink port to failover to a standby VC uplink port?

A21: Depending on the upstream networking technology used - it can be as long as 2 seconds or less for an optimized configuration

Q22: How fast should I expect a port channel (LAG) to failover? A22: 2 seconds or less for an optimized configuration.

Q23: How fast should I expect my teamed NICs to failover when a vNet fails over between uplinks?

A23: 2 seconds or less for an optimized configuration.

Q24: Can I connect VC-Enet's XFP ports to XENPACK or X2 ports on a Cisco switch?
A24: Cisco states that XENPACK, X2, and XFP modules are compatible if using the same port type. Since VC only supports 10GB-LR and 10GB-SR, Virtual Connect should be compatible with XENPACK, X2, and XFP modules from Cisco as long as they are using 10GBASE-LR or 10GBASE-SR modules. Since Cisco supports several port types, several transceiver types, and many different interface modules for their switches, HP does not test every combination.
http://www.cisco.com/en/US/prod/collateral/modules/ps5455/prod_brochure0900aecd8034bba6.pdf

Q25: Is Virtual Connect compatible with layer 3 protocols other than IP? For instance, does VC support IPv6,

IPX, AppleTalk, etc.?
A25: Virtual Connect only supports IP (IPv4) on its management interfaces (Web, SSH CLI, or SNMP). In reference to Virtual Connect's bridging functionality, VC supports any layer 3 or higher protocol in use on blade servers. Since Virtual Connect is a layer 2 device, it is layer 3 protocol agnostic. Meaning, the blade servers can communicate through VC using any upper layer protocol (e.g. IPv4, IPv6, IPX, AppleTalk, etc.) that's carried within an Ethernet frame.

Q26: Does Virtual Connect support jumbo frames?
A26: Yes, VC-Enet supports Ethernet frames sizes up to 9216 bytes.

Q27: Does Virtual Connect support EtherChannel\802.3ad\SLB on the downlinks to the server NICs? Can I use LACP port trunking on the server NICs connected to Virtual Connect?
A27: No, Virtual Connect does not support EtherChannel\802.3ad on the downlinks to server NIC ports.

Q28: What Level of SSH does Virtual Connect support?
A28: Virtual Connect has a strong security across management interfaces with support for SSL and SSH —including 2048-bit SSL certificates.

Q29: What is the inter-port latency.on the Virtual connect Ethernet modules.
A29: The following is the current numbers:

FlexFabric	1.2 µs on Ethernet only ports
	1.7 µs Ethernet/FC port
Flex 10/10D	.9 µs on Ethernet only ports
Flex 10	1.2 µs on Ethernet only ports

Q30: Can I access Virtual Connect from the CLI and what are the CLI commands.
A30: Virtual connect has a robust CLI command structure. The same commands that are executed in the GUI can be executed in the CLI. For a complete listing of commands seethe following link: Virtual Connect CLI Commands

© 2014 Hewlett-Packard Development Company, L.P. The information contained herein is subject to change without notice. The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

VMware and VMware ESX server is a trademark or registered trademark of VMware, INC. or its subsidiaries in the United States and other countries.

Cisco and EtherChannel are trademarks or registered trademarks of Cisco Systems, Inc. or its subsidiaries in the United States and other countries.

January 2014